



Neutral Citation Number: [2019] EWHC 1687 (Pat)

Case No: HP-2017-000048

IN THE HIGH COURT OF JUSTICE
BUSINESS AND PROPERTY COURTS
INTELLECTUAL PROPERTY LIST (CHANCERY DIVISION)
PATENTS COURT

Roll Building
Fetter Lane, London, EC4A 1NL

Date: 4 July 2019

Before :

MR JUSTICE ARNOLD

Between :

CONVERSANT WIRELESS LICENSING SARL

- and -

(1) HUAWEI TECHNOLOGIES CO., LIMITED

**(2) HUAWEI TECHNOLOGIES (UK) CO.,
LIMITED**

(3) ZTE CORPORATION

(4) ZTE (UK) LIMITED

Claimant

Defendants

Adrian Speck QC, Mark Chacksfield QC and Michael Conway (instructed by EIP Legal)
for the Claimant

**Daniel Alexander QC and Miles Copeland (instructed by Allen & Overy LLP) for the First
and Second Defendants and (instructed by Bristows LLP) for the Third and Fourth
Defendants**

Hearing dates: 14, 17-19, 21, 24 June 2019

Approved Judgment

I direct that pursuant to CPR PD 39A para 6.1 no official shorthand note shall be taken of this Judgment and that copies of this version as handed down may be treated as authentic.

.....
MR JUSTICE ARNOLD

MR JUSTICE ARNOLD :

Contents

<i>Topic</i>	<i>Paragraphs</i>
Introduction	1-2
The witnesses	3-11
Prof Marshall	3-7
Dr Brydon	8-11
The skilled person	12
Common general knowledge	13-65
UMTS	15-16
Basic network infrastructure in UMTS	17-18
CDMA	19-20
Duplexing	21
Layers, protocols and stacks	22-26
Transmission Time Interval	27-30
HSDPA	31-32
HSUPA	33-35
E-DCH architecture	36-39
E-TFCs	40
Rise over Thermal	41-44
Scheduled and non-scheduled grants	45-50
The power headroom	51
The minimum set	52-53
HARQ	54-61
E-TFC selection and E-DCH transmissions	62-65
The Patent	66-147
Field of the invention	67-68
Description of the related art	69-75
Summary of the invention	76-90
Detailed description of the presently preferred embodiments	91-147
The claims	148-150
Construction	151-167
[A] A method...for autonomous uplink transmission	152-155
[B] a virtual transmission time interval for a medium access control entity	156-159
[B] which virtual transmission time interval defines a minimum time interval that is allowed between enhanced uplink transmissions	160
[C] data packets	161
[C] checking to determine whether the medium access control entity is transmitting data packets in a current air interface transmission time interval, by checking whether the medium access control entity is able to empty the radio link control buffer in the current air interface transmission time interval	162-167
Essentiality	168-205
Uplink DTX	168-169
Uplink DRX	170-179

Buffer Occupancy	180-184
Retransmissions in Uplink DRX	185-190
Conversant’s case in outline	191-194
Integer A	191
Integer B	192
Integer C and D	193-194
Assessment	195-205
Integer A	195
Integer B	196
Integer C	197-205
Added matter	206-227
The law	206
Assessment	207-209
Allegation 1	210-213
Allegation 2	214-215
Allegation 3	217-218
Allegation 4	219-226
Allegation 5	227
Obviousness	228-258
Samsung	229-244
The disclosure	229-234
Obviousness	235-245
Terry	246-258
The disclosure	246-25
Obviousness	256-258
Insufficiency	259
Summary of principal conclusions	260

Introduction

1. The Claimant (“Conversant”) is the proprietor of European Patent (UK) No 1 797 659 entitled “Slow MAC-e for autonomous transmission in high speed uplink packet access (HSUPA) along with service specific transmission time control” (“the Patent”). Conversant contends that the Patent is essential to an aspect of the third generation mobile phone Universal Mobile Telecommunications Standard (“UMTS”) referred to as Uplink DRX and thus has been infringed by the Defendants, who market mobile phones which comply with UMTS. The Patent forms part of a portfolio owned by Conversant which is subject to a FRAND undertaking, but so far no licence has been agreed between Conversant and the Defendants. The Defendants dispute essentiality, and hence infringement, and counterclaim for revocation on the grounds of added matter, obviousness and insufficiency. Conversant has applied unconditionally to amend the claims of the Patent. It does not admit that the granted claims are invalid, but it does not contend that they are valid. The filing date of the Patent is 30 September 2005. Conversant does not maintain any claim to priority from two priority documents, the earliest of which was filed on 1 October 2004. As will appear, the specification of the Patent refers both to proposals which were current as at October 2004 and to later developments of UMTS.
2. It should be noted that the Defendants dispute the jurisdiction of the courts of England and Wales on the basis that Conversant seeks the determination of a global FRAND

licence, whereas the Defendants contend that the courts of England and Wales only have jurisdiction to determine a FRAND licence for the UK. Although the Defendants' jurisdictional challenge has not so far been successful, the Supreme Court has given the Defendants permission to appeal and the appeal is due to be heard in October 2019. The Defendants' participation in this trial was without prejudice to that challenge.

The witnesses

Prof Marshall

3. Conversant's expert was Professor Alan Marshall. Prof Marshall is a Professor of Communications Networks in the School of Electrical Engineering, Electronics and Computer Science at the University of Liverpool, head of the Department of Electrical Engineering and Electronics and Director of the Advanced Networks Research Group. He obtained a BSc in Microelectronics Systems from the University of Ulster in 1985 and a PhD on the development of a local area network from the University of Aberdeen in 1990. From September 1990 to April 1993 he was employed as a senior systems engineer by Northern Telecom (UK) and Bell Northern Research working on the second generation system GSM. From 1993-1997 he was Lecturer, from 1997-2001 Senior Lecturer and from 2001 to 2013 Professor of Telecommunications Engineering and Networks at the Queen's University, Belfast. He moved to Liverpool in 2013.
4. As well as his research into wireless networks since 1993, Prof Marshall has acted as consultant for, and given industrial training courses on, third generation wireless networks, including for Alcatel Lucent and Flextronics (formerly Nortel) in 2005-2006. He is a Fellow of the Institution of Engineering and Technology and a Senior Member of the Institute of Electrical and Electronic Engineers (IEEE).
5. Counsel for the Defendants submitted that Prof Marshall had less relevant expertise than Dr Brydon since he had not been involved in developing recent telecoms standards; but neither had Dr Brydon. Prof Marshall accepted in his third report that he had made a number of minor technical errors in his first report which Dr Brydon had pointed out, but none of these were of significance.
6. Counsel also submitted that, because of the way in which Prof Marshall had been instructed, his evidence was affected by hindsight. This is because Prof Marshall was shown the Patent as proposed to be amended, rather than the application for the Patent as filed ("the Application") and because he was aware of Uplink DRX when considering the disclosure of the Patent. I agree that this procedure had the potential to lead Prof Marshall to interpret the Patent with hindsight, but it does not necessarily follow that his interpretation was tainted by hindsight. In any event, there is no dispute that the interpretation of the Patent is a question for the Court, albeit guided by the expert evidence, nor that the Patent must be interpreted as at the filing date without hindsight.
7. Finally, counsel pointed out that aspects of Prof Marshall's oral evidence departed from his written evidence. In my view this is mainly indicative of Prof Marshall's willingness to concede points which were put to him in cross-examination. As counsel for Conversant submitted, however, there were occasions when the cross-examiner

and the witness were at cross-purposes, and therefore some caution is required in assessing the evidence.

Dr Brydon

8. The Defendants' expert was Dr Alastair Brydon. Dr Brydon received a BSc in Electronics from University of Manchester Institute of Science and Technology in 1984 and completed a doctorate in Medium Rate Data Transmission at High Frequency at UMIST in 1990. During his PhD, he also worked as a lecturer.
9. From 1989 to 1995 Dr Brydon was employed by BT Research Laboratories as a Senior Engineer, working on research and standardisation projects in respect of GSM, UMTS, DECT and TETS mobile radio systems. His work resulted in patents in Europe and the US. He chaired the co-ordination group in the European Community's MONET project, which brought together 28 European telecommunication companies to develop fundamental aspects of the network standards for UMTS. From 1995 to 1997 Dr Brydon worked for Cellnet as their Network Architecture Manager. From 1997 to 2001 he worked for Nokia Networks, where he was responsible for strategy and business development of mobile services in EMEA countries.
10. In 2001 Dr Brydon co-founded Sound Partners Ltd (subsequently renamed Unwired Insight Ltd), a research and consultancy business focussing on mobile technology and services. Unwired Insight has provided technical and commercial advice and research in respect of 2G to 5G (inclusive) mobile communication systems to network operators, equipment manufacturers, investors and regulatory bodies. Dr Brydon has authored more than 40 major reports and publications. He is a Chartered Engineer, a Fellow of the Institution of Engineering and Technology and a Senior Member of the IEEE.
11. Counsel for Conversant rightly accepted that Dr Brydon was a good witness and made no criticism of him.

The skilled person

12. It is common ground that the skilled person to whom the Patent is directed is a systems engineer with responsibility for layer 2 aspects of UMTS, and in particular the Medium Access Control ("MAC") sub-layer. Such a person would have a degree in electrical engineering or a similar subject and 3-5 years' experience in cellular communications.

Common general knowledge

13. It is common ground the skilled person would be familiar with the basic operation of UMTS and would be aware in general terms of the contents of the 3GPP technical specifications insofar as relevant to their particular area of expertise. That would include awareness of the latest versions of the MAC specification (TS 25.321 v6.5.0 June 2005) and the High Speed Uplink Packet Access (HSUPA) overall description (TS 25.309 v6.3.0 June 2005).
14. Regrettably the parties did not agree a technical primer in this case despite the encouragement to do so contained in what is now paragraph 14.6 of the Patents Court


Guide. This should not have been a problem given that there was little dispute between the experts as to the relevant technical background. Counsel for Conversant helpfully set out in their skeleton argument a summary of the technical background, certain aspects of which they expanded upon in their written closing submissions. Counsel for the Defendants took no point on the accuracy of that summary, nor did he dispute that all of it would have been part of the common general knowledge of the skilled person. Accordingly, I have largely based the following account on that summary.

UMTS

15. UMTS is specified by 3GPP. The Standard consists of a large number of Technical Specification (TS) documents describing different aspects of the network such as those referred to in paragraph 13 above. These evolve over time as new Releases of UMTS are introduced containing new features.
16. UMTS was first introduced in 1999 (Release '99). Prior to the filing date there had been several further releases. Release 4 was largely completed by 2001. Release 5, which was largely completed in 2002, introduced High Speed Downlink Packet Access (HSDPA, described below). By the filing date, Release 6, which defined HSUPA in an early form, had been introduced.

Basic network infrastructure in UMTS

17. Broadly, the UMTS network can be thought of as divided into three parts: the user equipment (UE) (generally a mobile phone); the Radio Access Network (RAN, also referred to as the UMTS RAN or UTRAN), which handles transmissions between multiple UEs and the Core Network; and the Core Network, which manages and processes transmissions across the network and interfaces with external networks such as the internet. This is shown schematically in the following diagram:



18. As shown, the RAN consists of the Node B (also called a base station) and the Radio Network Controller (RNC). Each Node B communicates with several UEs in an area referred to as a cell. Transmissions of data occur between the UE and the Node B

(called the uplink or UL) and in the reverse direction between the Node B and the UE (called the downlink or DL). A single RNC controls and coordinates several Node Bs.

CDMA

19. Because several UEs may be in communication with a single Node B, it is necessary to have a mechanism for the Node B to distinguish the transmissions of different UEs in the uplink and for a UE to determine which transmissions are intended for it in the downlink. The way in which radio resource is shared among multiple users is referred to as “multiple access”. UMTS uses a method called Code Division Multiple Access (CDMA), also referred to as Wideband CDMA or WCDMA, in which transmissions are sent at the same time and using the same frequency, but can be distinguished through the use of different spreading codes.
20. An important consequence of the use of a code division system is that the transmissions of other users are perceived as noise at the Node B receiver in relation to the desired signal – i.e. the signals interfere with each other. This leads to the system being “noise-limited”: the greater the number of users transmitting at the same time, the higher the level of interference experienced in relation to the desired signal.

Duplexing

21. Duplexing refers to the separation of uplink and downlink transmissions. Two methods are used:
 - i) Frequency division duplexing (FDD), which allocates separate frequency bands to the uplink and downlink. Transmissions may then be sent simultaneously on the uplink and downlink.
 - ii) Time division duplexing (TDD), which enables uplink and downlink communications by time-division of transmissions on the same frequency band. In TDD, specific time slots are allocated for the UE to receive bursts of data on the downlink, and then to transmit a burst of data on the uplink.

Layers, protocols and stacks

22. The Open Systems Interconnection (OSI) model is a conceptual model which is used to explain functional layers in a communications network. It consists of seven layers.
23. Each layer transmits data on behalf of the layer above, and in turn relies on the services of the layer below to transmit its output to the corresponding layer in the receiver. Data packets arriving from the above are called Service Data Units (SDUs). Once a SDU has been processed by a layer, it is then passed to the layer below as a Protocol Data Unit (PDU). Where needed, the layer will add its own control information header, containing for example a destination address.
24. Each layer is considered to communicate via a logical link with its peer layer in the transmitter or receiver.
25. Data transmissions in UMTS are organised on the basis of a protocol layer stack. Each layer is defined by a shared communications convention or “protocol”, so entities within the same layer can communicate with each other, but do not have to

know how entities in a different layer of the stack communicate with their corresponding partner entities. The different layers can therefore be ignorant of each other, save for the protocols necessary to pass information and queries up and down the stack.

26. Within UMTS, the different layers which are relevant to this case can be characterised as follows:
- i) Layer 1 is the physical layer (PHY), which is responsible for transmitting and receiving data over the airwaves. It maps transport channel formats specified by layer 2 onto physical channel formats which are the formats actually sent over the air. It is in the PHY that the processes of spreading and modulation are carried out.
 - ii) Layer 2 is the data link layer. It ensures there is a reliable flow of information across a communications link e.g. a radio link in UMTS. It consists of three sub-layers: the MAC, the Radio Link Control (RLC), which sits above the MAC; and the Packet Data Convergence Protocol (PDCP), which sits above the RLC. The RLC sends data to, and receives data from, higher layers and segments into, or reassembles it from, smaller RLC PDUs, of defined length. The RLC also buffers data for transmission. The MAC is responsible for mapping data from different services or logical channels onto transport channels for transmission by the PHY, and for multiplexing or demultiplexing higher layer PDUs into/from transport blocks to be sent to or received from the PHY.
 - iii) Layer 3 is the network layer. The function of this layer is to route packets across different networks.

Transmission Time Interval

27. In UMTS transmission periods at the PHY are divided into 10 ms radio “frames”. Each frame is sub-divided into 15 “timeslots”, with three timeslots forming a “sub-frame” of 2 ms. There are therefore five sub-frames within a frame.
28. When considering the PHY, the Transmission Time Interval (TTI) defines the duration of time over which the PHY may transmit a transport block over the air. In versions of UMTS prior to Release 5, the TTI could be of 10 ms, 20 ms, 40 ms or 80 ms. Release 5 introduced a high speed downlink channel (HSDPA, discussed below) into which a shorter 2 ms TTI was introduced – equivalent to 1 sub-frame. Later, Release 6 introduced a high speed uplink channel (HSUPA, also discussed below) with a 2 ms TTI, although a 10 ms TTI was retained.
29. The TTI for transmission over the air interface has strictly defined start and end points, and is of fixed duration. The processing steps conducted by each layer in relation to the data that will be transmitted over the air interface during the course of a given TTI cannot be performed simultaneously. Each layer must finish its processing before the next layer down is able to do so. This may also involve necessary steps both prior to and after transmitting data to the next layer in the protocol stack. When discussing a TTI, the context is the period of time in relation to which the particular layer under discussion is performing the necessary processing relating to the data

which will ultimately be sent out over the air interface in that TTI. A TTI in different layers cannot be equated with some particular instant of time on a clock.

30. “Transmitting” and “transmission” generally refers to the process of handing down data from one layer in the protocol stack to the next. It is only at the PHY that “transmitting” equates to transmission over the air interface.

HSDPA

31. The goal of HSDPA was to allow higher speed transmissions on the downlink to meet an increased demand for higher data-rate services such as web-browsing.
32. HSDPA introduced several important changes. Of particular note were the following:
 - i) The introduction of the 2 ms TTI (as mentioned above).
 - ii) The introduction of fast scheduling at the Node B. This transferred scheduling from the RNC to the Node B, allowing scheduling to be performed more quickly and responsively to the conditions on the air interface.
 - iii) The use of adaptive modulation and coding (AMC). This is a way of dynamically altering the modulation and/or coding used on the downlink to account for variation in signal quality. Importantly, the power of the downlink signal remains constant in respect of each UE (something that cannot be achieved on the uplink).
 - iv) Multi-code operation enabled multiple streams of data to be sent to a single UE or between multiple users each TTI by the use of spreading codes.
 - v) The introduction of HARQ retransmissions (as to which, see below).

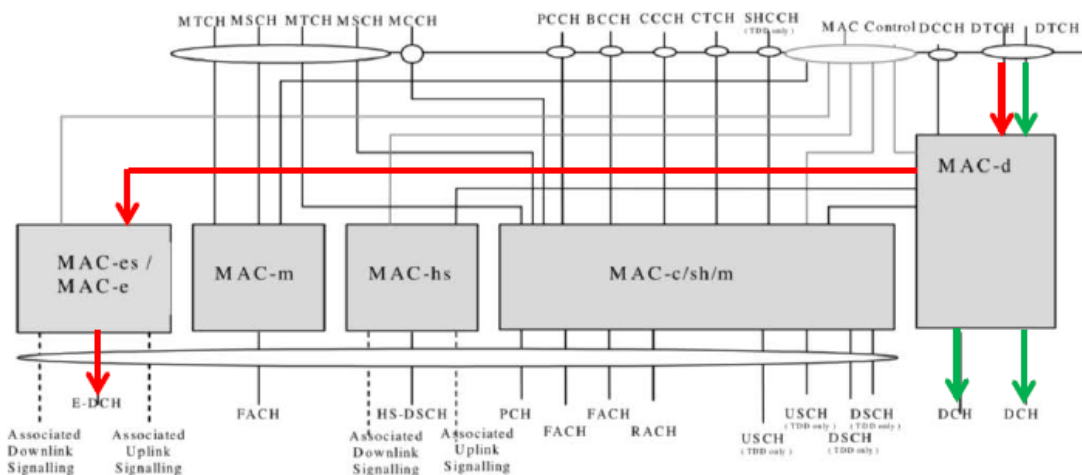
HSUPA

33. As noted above, HSUPA was introduced in Release 6. Release 6 was developed to completion over a period ending very shortly before the filing date.
34. The basic premise of HSUPA was to enable high speed transmissions in the uplink alongside the high speed downlink that had been introduced in Release 5. This was implemented through a new dedicated uplink channel, called the Enhanced Dedicated Channel (E-DCH, described below). The concept of the uplink E-DCH was first introduced in January 2005.
35. In developing HSUPA, the developments that had been introduced to HSDPA in relation to the downlink could not simply be replicated in the uplink. This is due to some fundamental differences between the uplink and downlink, including:
 - i) The downlink is a one-to-many connection whereas the uplink is a many-to-one. This means that in the downlink power resource is centralised, whereas in the uplink it is not feasible to have transmissions at a constant power. The reason the 10 ms TTI was retained in HSUPA was to allow transmissions at lower power from power-limited UEs on the cell edge.

- ii) In the downlink, transmissions are coordinated at the Node B, which is aware of the data that needs to be buffered and transmitted to each UE in the cell. In the uplink, transmissions are not coordinated: the network is not aware of the details of the data that each UE is generating, their power limitations etc, and none of the UEs are aware of each other at all.
- iii) In the downlink, the transmissions to each UE are precisely synchronised, and hence the codes assigned to each UE are orthogonal to each other. This is not the case in the uplink, and therefore interference between the transmissions of different users is more problematic in the uplink (as explained below).

E-DCH architecture

- 36. Prior to the introduction of the E-DCH, MAC processing that was carried out for dedicated traffic was done by a component called MAC-d (“d” for dedicated).
- 37. The introduction of E-DCH involved significant changes in the MAC architecture at both the UE and the Node B. The key change was the introduction of a MAC-e component at the Node B and a MAC-es component in the RNC, and of a combined MAC-e and MAC-es component at the UE (“e” for enhanced). These components were introduced to manage the new functionality associated with the E-DCH, such as scheduling and HARQ processing. In the UE the MAC-e/es component sits underneath MAC-d and is fed by it. MAC-es receives PDUs from particular logical channels directly from MAC-d, and combines these into MAC-es PDUs. Several MAC-es PDUs may then be combined at the MAC-e layer into a single MAC-e PDU.
- 38. The structure of the new MAC layer in the UE is shown in the following diagram (based on figure 4.2.3.1 in TS 25.321 v6.5.0 June 2005):



- 39. It can be seen from this diagram that the E-DCH is one part of the data that flows through MAC-d. For the E-DCH data (the red line) the functionality for controlling transmissions was transferred to the MAC-e. Thus, while a number of new parameters relating to the control of E-DCH transmissions was introduced into the MAC-d, the functionality which acted upon them was implemented in the MAC-e. In relation to enhanced traffic, the MAC-d is effectively transparent. The MAC-d continues to manage the legacy dedicated channels (the green lines) as before.

E-TFCs

40. E-TFCs (E-DCH Transport Format Combinations) are combinations of parameters permitted by UMTS from which the MAC can choose for the purpose of creating, and instructing the PHY to send, a MAC-e PDU. Each E-TFC comprises a combination of the transport block size (number of bits), TTI length, coding, and cyclic redundancy check which in turn correspond to different data rates. Generally, the higher the data rate, the more power that is required for the E-DCH transmission.

Rise over Thermal

41. Rise over Thermal (RoT) is a measure of interference within the cell that indicates the ratio between the total interference received on a base station and the thermal noise. RoT is a particular issue on the uplink in CDMA systems, where a number of UEs transmit to the same Node B receiver and hence cause interference for each other.
42. As the number of UEs active in a cell increases, the level of interference associated with each individual signal also increases. The degree of interference depends not just on the number of UEs, but also on the power with which they are each transmitting (i.e. how much data). An increase in interference in a cell may also affect its neighbouring cells. As such, the system needs to manage that interference to provide an adequate service within cells and to avoid unnecessary interference in neighbouring cells.
43. The requirement to control the overall level of noise within the cell in this way gives rise to the need to control the power used by each UE to transmit data on the uplink. This is achieved by resource allocation by the RNC, and dynamic scheduling (high speed control over transmission power by individual UEs) by the Node B, aiming to operate the cell close to its maximum planned load, so it carries as much traffic as possible, but without straying into the overload area.
44. In HSUPA, the Node B is responsible for allocating a power budget to UEs on a cell-wide and TTI-by-TTI basis for data transmitted on the E-DCH (called scheduled transmissions – described in more detail below). Since Node B controlled transmissions are not the only kind of transmission, the Node B cannot know precisely which UEs will be transmitting at a particular time and with what power they are transmitting. Accordingly, the Node B must calculate an overall load within the cell, based on the number of UEs within the cell, and an estimated power usage by transmissions that are not under the Node B's control (i.e. non-E-DCH transmissions, which are prioritised, and non-scheduled E-DCH transmissions), as well as interference contributed by UEs outside the cell. Subtracting the estimate for these other categories from the maximum planned load allows the Node B to budget for what resources are available for it to allocate for scheduled transmissions. Decreasing this margin (which may be wasted if there is in fact no transmission outside the Node B's control) would allow the Node B to increase the budget for scheduled transmissions and hence improve overall efficiency in the system.

Scheduled and non-scheduled grants

45. The transmission of all data on the E-DCH requires a scheduled or a non-scheduled grant. The use of scheduled and non-scheduled grants allows a combination of fast

power control by the Node B (TTI by TTI) for some types of data with transmissions of other types of data within a power limit predetermined by the RNC.

46. The scheduled grant (sometimes called the Serving Grant) is the means by which the Node B exerts fast control over the available radio resources allocated to each UE for certain types of data. It specifies the power that the UE may use for transmissions across all scheduled data flows.
47. The Node B communicates the scheduled grant via two control channels, E-AGCH and E-RGCH. E-AGCH communicates the absolute grant, which specifies a value for the Serving Grant maintained by the UE, while E-RGCH communicates the Relative Grant, which is used to increment the Serving Grant each TTI (possible values being “up”, “down” or “hold”). The UEs can also send scheduling information on the uplink as part of an E-DCH transmission, by which the UE indicates its current status, including information on buffer status, traffic priority and power availability.
48. The non-scheduled grant is the means by which the resources predetermined by the RNC are signalled. These apply to data in non-scheduled flows. A non-scheduled grant is determined and signalled from the RNC, rather than the Node B, meaning that it cannot be controlled or varied as quickly as the scheduled grant. Data belonging to a non-scheduled flow is therefore sent without a grant from the Node B. The non-scheduled grant is used in particular for delay sensitive services, such as control transmissions, or Voice over IP (VoIP, a way of sending voice signals as a series of data packets).
49. Only data that has been allocated a non-scheduled grant may be sent using that non-scheduled grant, while only scheduled data may be sent using the scheduled grant – they are not interchangeable. The details of how the MAC-e PDU is constructed from scheduled and non-scheduled data are summarised below.
50. Accordingly, the sum of the scheduled and non-scheduled grants provides an upper limit for the data which may be transmitted in a particular TTI across both scheduled and non-scheduled data.

The power headroom

51. The MAC-e also has to determine what capacity it has to send data in a given TTI based on an additional limitation known as the “power headroom”. The power headroom is calculated on the basis of an absolute power limit for all uplink transmissions by the UE communicated by the RNC, from which the MAC-e subtracts the power required to transmit a number of other channels to calculate the power available for transmissions on the E-DCH. Since larger E-TFCs require higher power levels, the power headroom can be converted by the UE into a list of available E-TFCs from which the MAC-e is able to select in this TTI – those requiring too much power being excluded. The overall effect is that the MAC is limited in essence by the lower of the limits from either the sum of the scheduled plus non-scheduled grants, or the power headroom.

The minimum set

52. As at the earliest priority date of the Patent, when the discussions about the development of HSUPA were still at a relatively early stage, it had been proposed that the “minimum set” could provide a means by which data could be transmitted without a grant from the Node B (there being no concept at that time of a separate non-scheduled grant). By the filing date, however, the non-scheduled grant had been introduced for this purpose and the role of the minimum set was different.
53. The “minimum set” defines a set of E-TFCs whose use is always supported irrespective of the limitations of the power headroom. They are therefore always available to be used to carry scheduled and/or non-scheduled data, although that data still has to be the subject of a valid grant. Thus the minimum set still provides an exception, but to a different rule.

HARQ

54. Hybrid Automatic Repeat Request (“HARQ”) is a mechanism for combining forward error correction and automatic repeat requests. For present purposes, it is not necessary to describe HARQ in detail, but it is important to note the following matters.
55. HARQ allows data to be re-transmitted when it is not correctly or completely received by the Node B.
56. Each MAC-e PDU is assigned to a specific HARQ process, which handles the HARQ protocol that will be applied to transmissions and retransmissions of the MAC-e PDU associated with it.
57. The Standard specifies the number of HARQ processes available for E-DCH transmissions. For a 2 ms TTI there are eight HARQ processes; for a 10 ms TTI there are four. The HARQ processes are made available sequentially and in order, one per TTI. Each MAC-e PDU is assigned to a specific HARQ process such that it may only be transmitted (and, if necessary, retransmitted – see below) on that process.
58. The UE and the Node B therefore know which HARQ process is available to be used in each TTI, meaning that there is no need to send control information indicating the HARQ process being used for any particular transmission. This saves on control data, and is known as “synchronous HARQ” (as the UE and Node B are synchronised in this aspect).
59. For each uplink transmission, the Node B sends the UE an acknowledgement (ACK) or negative acknowledgement (NACK) signal on a dedicated channel called the E-DCH HARQ Indicator Channel (E-HICH) to inform the UE whether that transmission has been correctly received.
60. If the transmission is not successfully confirmed, then the UE will repeat the transmission. This is known as a retransmission. Retransmissions may only be sent in the same HARQ process as the original transmission. Retransmissions are interleaved with ongoing new transmissions. Only a certain number of retransmission attempts are allowed. When an ACK is received, or the permitted number of retransmission attempts is exhausted, the data that was occupying that HARQ process is discarded and a new transmission may be made.

61. As a result, until a particular packet has been ACKed by the receiver, the relevant HARQ process will remain blocked for any other transmissions to be sent, until the maximum limit on the number of HARQ retransmissions has been exhausted. With, for example, a 10 ms TTI, four HARQ processes and four HARQ retransmissions, a HARQ process may remain blocked for 160 ms by a single MAC-e PDU.

E-TFC selection and E-DCH transmissions

62. Before any new E-DCH transmission, E-TFC selection must be performed. This is the process by which the MAC-e allocates data and determines what transport format will be used in order to construct a MAC-e PDU.
63. As explained above, generally a MAC-e PDU will comprise both scheduled and non-scheduled data. To briefly summarise:
- i) The MAC-e determines the scheduled and non-scheduled grants, and power headroom. These values feed into the process of multiplexing and E-TFC selection.
 - ii) The power headroom is used to determine the set of E-TFCs from which E-TFC selection may take place.
 - iii) The MAC-e identifies those flows for which there is available data and allocates that data, starting with the highest priority data. If there is no data for the highest priority flow, the MAC-e will consider the next highest priority, and so on. In this way the MAC-e allocates data up to the limits defined by the non-scheduled grant, scheduled grant and the power headroom.
 - iv) Based on the amount of data available, the MAC-e selects an available E-TFC format. The standards specify that the MAC-e will select the E-TFC format that will result in the minimum amount of padding.
 - v) The “multiplexing and TSN setting” entity then constructs a MAC-e PDU. One or several MAC-d PDUs from a particular MAC-d flow may be multiplexed into a MAC-es PDU, to which a MAC-es header is added. Several MAC-es PDUs and their headers may then be concatenated into a MAC-e PDU.
64. Prof Marshall illustrated the process by way of a worked example, adapted from a textbook, in Figure 12 of his first report:

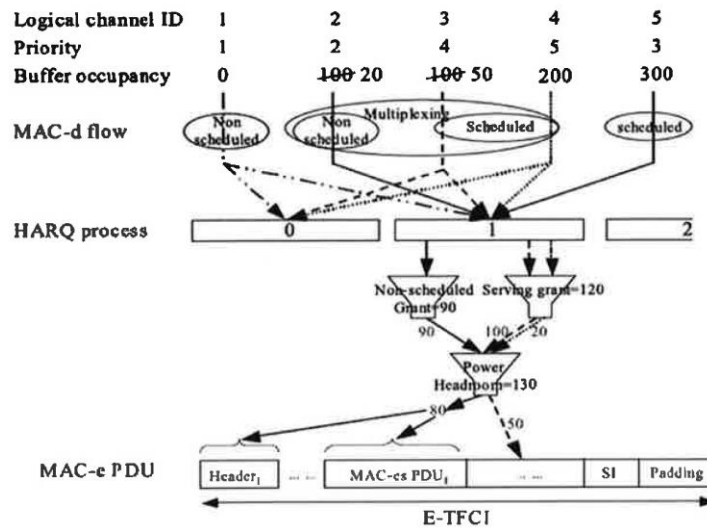


Figure 12 - Worked example of bit mapping for E-TFC selection (Reproduced from L. Hanzo, J.S. Blough. S. NI (2008) "3G, HSPA and FDD vs TDD Networking" ISBN 978-0-470-75420-7 (HB) Chapter 2, Section 2.3.2.

65. The example shows a process of E-TFC selection, multiplexing and MAC-e PDU generation with five logical channels. Channels 1 and 2 are non-scheduled, and channels 3-5 are scheduled (with channels 3 and 4 being mapped onto the same MAC-d flow). As shown in the figure, the process of allocating data from the MAC-d flows requires the MAC-e to have regard to the available data in each MAC-d flow, its priority, multiplexing ability and the level of the scheduled and non-scheduled grants.

The Patent

66. A large part of the dispute in this case concerns the disclosure of the Patent. Although the specification is not very long, it is quite densely packed. Understanding is not assisted by the fact that it contains some errors. As the disagreements between the experts show, parts of it are capable of being read in more than one way. Furthermore, it is the Defendants' case that Conversant's case amounts to an attempt to reconstruct the claimed invention with the benefit of hindsight so as to make it read on to Uplink DRX. Conversant not only disputes this, but points out that, despite that being their case, the Defendants did not in their submissions analyse the disclosure of the Patent separately from their arguments on essentiality. It is therefore important to analyse the disclosure with care before turning to the case on essentiality. I shall do so under the headings in the specification.

Field of the invention

67. The specification states at [0001] that the invention relates to the UMTS RAN standards, HSUPA and HSDPA, and more particularly to "a system and methods for slow ... MAC-e for autonomous transmission during HSUPA, an[d] for service specific transmission time control in HSUPA". It is common ground that the skilled reader would understand from this that, consistently with the title of the Patent, the invention has two aspects, one relating to slow MAC-e for autonomous transmission and one relating to service specific transmission time control. As the skilled reader

will discover, the typographical error “an” is the first of a number of errors in the specification (not all of which it is necessary for me to mention).

68. The expression “autonomous transmission” is not defined in the Patent, but there is no dispute that the skilled reader would understand from reading the specification that an autonomous transmission is one for which a valid scheduled grant is not required from the Node B in order for the packets to be transmitted.

Description of the related art

69. In [0002] the specification notes that TS 25.309 established that some level of E-DCH minimum set support is required to provide backwards compatibility. It goes on to say that “under normal conditions” the UE is only able to transmit if it has received a scheduling grant from the Node B, but the minimum set provides a set of transport formats for which a scheduling grant is not required. This reflects the earlier proposal in which the minimum set was proposed to be used when no Node B resources had been allocated to that UE.
70. At [0003] and [0004] the specification discusses two problems with autonomous transmissions. The first problem is that, because a Node B must be continuously ready to process unscheduled autonomous transmissions irrespective of the number of UEs that actually perform an autonomous transmission, there is a limit to the ability of the Node B to optimise its processing resources by scheduling.
71. The second problem concerns the need to reserve an RoT margin for UEs performing autonomous transmissions. Although not every UE will actually transmit in any particular time slot, the Node B cannot predict which UE(s) will transmit or when. The reservation of a RoT margin is an approach based on providing a margin which should be at a high enough level to be exceeded only (acceptably) rarely, but as low as possible to make most efficient use of the available resources. As a consequence, however, the network needs to reserve a significant margin for autonomous transmissions. Consequently, this reserved margin will reduce the amount of power available for scheduled transmissions. The specification notes that, the more UEs potentially able to perform autonomous transmissions, the greater the margin that the Node B must reserve, which “would degrade the performance of the scheduled transmissions”. It is common ground that this latter statement would be understood to mean that a larger reserved margin would limit the amount of transmission resource that can be allocated to scheduled transmissions.
72. In [0005] the specification discusses three prior art proposals to address these problems (including Samsung, discussed below). It says that these proposals all require a high level of complexity. Prof Marshall accepted that all three proposals were concerned specifically with autonomous transmissions.
73. At [0006] to [0008] the specification mentions certain features of the later version of UMTS Release 6. In brief summary, [0006] introduces E-DCH, [0007] introduces the limitations on the HARQ processes that may be applied in E-DCH transmissions, and [0008] sets out that TTIs can be either 2 ms or 10 ms. These paragraphs distinguish repeatedly between scheduled and non-scheduled grants and between scheduled and non-scheduled MAC-d flows.

74. In [0009] the specification introduces another issue: the transmission of a low bit-rate service (i.e. one where there will be transmissions of small amounts of data) over E-DCH can lead to a large control overhead, particularly if transmissions are frequent, because the amount of control bits per TTI is the same for all packet sizes. It also explains that each transmission will require subsequent ACKs and NACKs in the downlink. It says that it is possible to reduce the control overhead “by transmitting more packets in the same transport block but less often” (i.e. grouping them together).
75. At [0010] the specification says that it would be preferable to “increase the transmission interval for specific services” such as VoIP in the UTRAN “in order to increase transmission capacity”. The UTRAN should take the characteristics of the service (e.g. bit rate and delay requirement) into account when defining the transmission interval for that service. [0011] expands upon this by reference to the example of VoIP. The specification points out that VoIP generates data packets every 20 ms, but speech can tolerate a degree of delay. Therefore it proposes sending data packets less frequently (once every 40 ms or 60 ms) in order to improve transmission capacity.

Summary of the invention

76. In [0013] the specification states that the invention relates to enhancing the uplink E-DCH for packet data traffic during HSUPA, and refers to two specific 3GPP documents.
77. At [0014] the specification introduces the concept of a control parameter that is independent of the air interface TTI, HARQ processes or E-DCH scheduling. It states:
- “This control defines a minimum time interval between subsequent new transmissions. The control has no impact on retransmissions, which are performed normally.”

It is common ground that “subsequent” here means successive.

78. It is also common ground that the specification distinguishes between new transmissions on the one hand and retransmissions on the other. It can be seen that it states that the control parameter does not affect retransmissions. As Dr Brydon explained, however, the skilled reader would understand from later parts of the specification (see in particular Figures 3 and 5 discussed below) that retransmissions may affect the timing of new transmissions, because a retransmission causes an increase in the time between a new transmission and the previous new transmission and a decrease in the time between a new transmission and the following one. As Dr Brydon accepted, the skilled reader would therefore appreciate that the idea of a minimum time interval between transmissions was subject to that effect.
79. The specification describes various approaches involving the application of this concept at [0015] to [0019]. These descriptions are repeated in the detailed description with further detail. At this stage, therefore, I will largely defer comment.
80. In a method described in [0015], for every MAC-e PDU, the UE checks for autonomous transmissions. If the transmission is not autonomous, the check is continually performed in a loop until there is an autonomous transmission. If an

autonomous transmission is detected, the system slows down (decelerates) the exchange rate between the MAC-e and the PHY i.e. the transmission of MAC-e PDUs between the MAC-e and the PHY. MAC-e PDUs are sent to the PHY only once every $n \times \text{TTI}$, instead of once every TTI.

81. At [0016] the specification says that, “[a]lternatively”, the rate at which the MAC-e sends PDUs to the PHY using transport formats belonging to the minimum set is decelerated so as to reduce the impact of the minimum set over RoT.
82. As discussed above, the skilled person would be aware that at an earlier stage it had been proposed to use the minimum set for sending transmissions autonomously. It is common ground that both these paragraphs would be understood as addressing that earlier proposal. It is also common ground that both paragraphs envisage implementation in the MAC-e layer and that this would not be service specific. This is, however, a dispute as to the disclosure of [0015] and [0016] which is more convenient to consider when I come to Figure 2.
83. The specification goes on in [0017] to say that, “[a]s the value of n is increased, the effective bit rate and the impact of the autonomous transmissions on RoT becomes smaller”. After a discussion concerning the HARQ processes, and the selection of n , the specification then states that the method is “advantageously simpler than conventional methods for performing scheduling grants for scheduled transmissions”, and that it is transparent to layer one and provides a minimum impact on layer two.
84. At [0018] the specification describes an embodiment of the invention in which a new control parameter is implemented either in a packet data protocol (PDP)/radio access bearer (RAB) layer or “in the MAC layer”. As Conversant points out, the reference here is to the MAC layer generally.
85. The former case, in which the control parameter is a new PDP context/quality of service (QoS) parameter, is discussed in the remainder of [0018]. The parameter is, in the preferred embodiment, a “service data unit (SDU) inter-arrival rate”, which establishes “the minimum time interval between consecutive SDUs transmitted on a specific RAB”. The specification states that:

“... the application does not deliver SDUs to the MAC layer at a higher rate than that specified by the parameter. If the data source produces several packets in this time interval, the packets are grouped into a single SDU”.

The skilled reader would understand from this that RLC SDUs are slowed down somewhere above the MAC layer, therefore affecting the subsequent flow of MAC-e PDUs into the PHY. It is then said that grouping the packets in the MAC layer enables the packet headers to be optimised and so more efficiently to share the uplink power resource between users.

86. The specification goes on at [0019]:

“Alternatively, it is possible to introduce a new MAC parameter in the MAC-d layer. In the preferred embodiment, the new parameter is a ‘virtual TTI’ that defines the minimum

time interval between subsequent new transmissions for a MAC-d flow. A first transmission would be permitted only once, during the virtual TTI. The virtual TTI could be signaled to the UE by the radio network controller (RNC). The UE could then implement the virtual TTI in the MAC-d layer.”

87. A number of points arise from this. First, it can be seen that the virtual TTI (“vTTI”) is introduced as a parameter which defines “the minimum time interval between subsequent [i.e. successive] new transmissions”. It is common ground that, although the point is not repeated here, it is clear from the end of [0014] and the detailed description later on in the specification that the parameter has no effect on retransmissions, which occur normally.
88. Secondly, as Dr Brydon agreed, the embodiment of [0019] is essentially the same as that in [0018] save that it operates at a lower layer. Due to the level at which the [0018] method is applied, however, this would be capable of being service specific; whereas the [0019] method could be MAC-d flow specific, but not service specific since a MAC-d flow can carry more than one service.
89. Thirdly, although [0019] refers to “a” MAC-d flow, Conversant contends that the language is equally apt to describe the situation in which multiple flows are affected; that there are technical reasons why the skilled person would not read it narrowly; and that a narrow reading is inconsistent with other parts of the disclosure e.g. [0043] (as to which, see below). I accept this.
90. Fourthly, the text refers to it being “possible” to introduce the new parameter in the MAC-d layer. As Conversant points out, this language is not mandatory. In any event, it does not limit where in the MAC the *functionality* acting upon that parameter could be located. The experts were agreed that the skilled reader would be aware of the distinction between where a parameter was stored, and the location of the functional software which made use of it, and in particular would be aware of a number of E-DCH parameters which were stored in the MAC-d, but the software implementing the functionality relating to these was in the MAC-e.

Detailed description of the presently preferred embodiments

91. The detailed description contains passages substantially repeated from the summary of the invention with further passages of detail and explanation added, in particular by reference to the figures. It starts at [0022] by repeating the end of [0001] and all of [0014]. This is followed by an explanation of the context of the invention by reference to Figure 1, but it is common ground that nothing turns on this.
92. Figure 2 illustrates the embodiment described in [0015], which is repeated with further detail at [0032]:

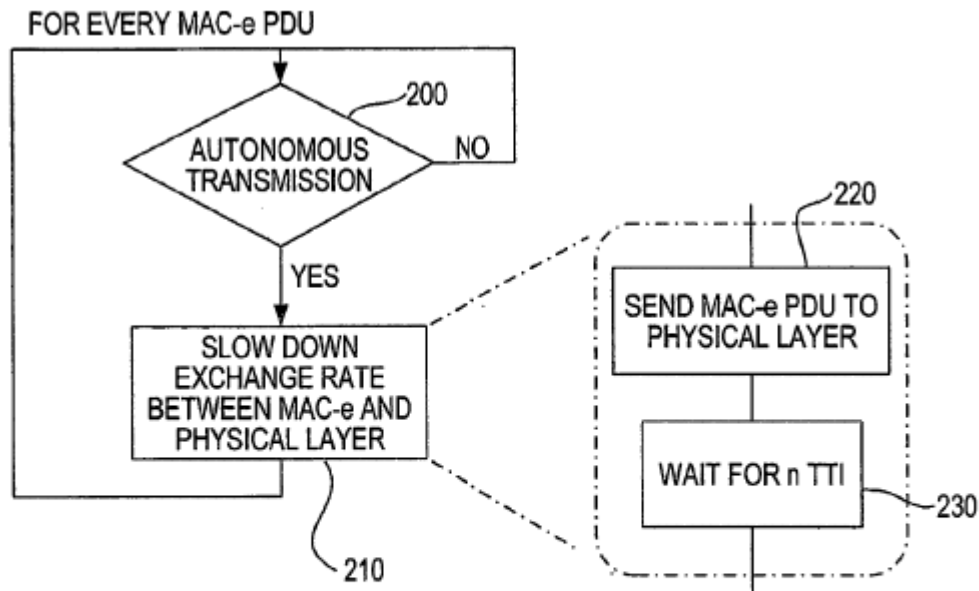


FIG. 2


93. As the specification explains at [0032]:

“For every MAC-e PDU, a check is made to determine whether the transmission is autonomous, as indicated in step 200. If the transmission is not autonomous, a check is continually performed until an autonomous transmission occurs, that is the method continues in a loop. If an autonomous transmission is detected, then the exchange rate between the MAC-e and the physical layer (layer one) is slowed down, i.e. the exchange rate is decelerated, as indicated in step 210. In accordance with the invention, the slow down of the exchange rate between the MAC-e and the physical layer occurs when the MAC-e layer ... sends a MAC-e PDU to the layer one (i.e., the physical layer), as indicated in step 220. The MAC-e PDU is sent to the physical layer every $n \cdot \text{TTI}$, instead of once every transmission time interval (TTI), as indicated in step 230.”

94. There is no dispute that the check is made for every MAC-e PDU whether it is an autonomous transmission or not, but there is a major dispute as to what happens if an autonomous transmission is detected. Conversant contends that, once an autonomous transmission is detected, the deceleration applies to all the following MAC-e PDUs. The Defendants contend that only autonomous MAC-e PDUs are slowed down, and thus the method disclosed is one that it is selective for autonomous transmissions.

95. In my view, the language of [0015] and [0032] is more consistent with Conversant's interpretation than the Defendants', since they speak generally of "the exchange rate between the MAC-e and the physical layer" being decelerated (not the exchange rate of autonomous transmissions) and since they say that "the slow down of the exchange rate ... occurs when the MAC-e layer ... sends a MAC-e PDU" to the PHY (i.e. any MAC-e PDU and not just an autonomous one). Furthermore, Figure 2 shows "MAC-e PDU" (i.e. any MAC-e PDU) being sent to the PHY at 220 followed by a wait for n TTI at 230 before the system goes back round the loop to continue the checking for an autonomous transmission.
96. Turning to Prof Marshall's evidence, his opinion was that this was the only way in which the specification could be interpreted since anything else arriving during the n TTI wait would be queued and since, otherwise, there would be different n TTIs for different autonomous transmissions. Prof Marshall did not accept that this reading would be contrary to the purpose of the invention, which was to address the problems identified at [0003]-[0004].
97. As for Dr Brydon, at one point in his evidence, he appeared to be saying that Figure 2 was unclear, consistently with the statement in his first report that the text of the specification was unclear and could be read as suggesting that all MAC-e PDUs were slowed down. Because Figure 2 did not show what happened to scheduled transmissions, his view was that the specification was "not necessarily saying anything about scheduled transmissions". While it was clear that autonomous transmissions were slowed down, Figure 2 "did not shout out to me as slow down scheduled transmissions as well". Later, he said that his interpretation was that the dotted box in Figure 2 was just referring to the autonomous transmissions, but there is nothing in the Figure to indicate that.
98. Turning to the alternative described in [0016], this is repeated almost verbatim at [0033]. There was some confusion concerning this alternative, but in the end the experts were largely agreed. In his first report, Dr Brydon suggested that [0016] was describing the same thing as [0015] in different words. Prof Marshall disagreed with this in his second report, suggesting that the methods were overlapping but different in that the method in [0016] involved checking for a minimum set. As he explained in cross-examination, the minimum set could encompass both autonomous and scheduled transmissions. Dr Brydon accepted that it was possible to interpret [0016] in the manner suggested by Prof Marshall. As Dr Brydon explained, on this interpretation, the alternative being suggested in [0016] is to check for the minimum set, which might be autonomous or scheduled, and slow those down. As Dr Brydon accepted, this would involve slowing down some scheduled transmissions as well as autonomous transmissions.
99. The upshot is as follows. So far as [0015], [0032] and Figure 2 are concerned, I prefer Conversant's interpretation to the Defendants' interpretation for the reasons given above. In any event, however, the conclusion I draw from the expert evidence is that the alternative described in [0016] and [0033] would involve slowing down some scheduled transmissions as well as autonomous transmissions.
100. [0034] repeats [0017] verbatim except for the addition of a reference to the definition of synchronous HARQ in TS 25.808.

101. [0035] repeats [0018] almost verbatim. Further detail is provided as to this embodiment at [0036] to the start of [0038] by reference to Figures 3 and 4.
102. The specification explains in [0037] that Figures 3 and 4 are exemplary timing diagrams showing the effect of the introduction of the PDP/RAB QoS parameter with an SDU inter-arrival rate of 40 ms in respect of, respectively, 10 ms and 2 ms TTIs. (It can be seen from the figures that they also involve four and eight HARQ processes respectively.) The diagrams do not show a specific layer. Rather they illustrate aspects of the system from different layers. The focus of these diagrams is to show the overall timing from the point of data generation by a VoIP application and its eventual successful reception at the receiving end.
103. When considering these diagrams, and Figures 5-7, it should be noted that the specification frequently refers to RLC SDUs (or just “SDUs”), that is to say, the inputs into the RLC layer, which lies above the MAC. The discussion tracks these through into the MAC layer, but by then they will have additional headers and control information associated with them, and the MAC PDUs may contain other data as well.
104. It is sufficient for present purposes to consider Figure 3:



105. The specification explains at [0037]:

“With reference to FIG. 3, for a 10 ms air interface TTI of 10 ms, 2 VoIP packets (i.e., 1 RLC SDU) are transmitted every 40ms. In this case, a single VoIP packet is transmitted from the source every 20 ms. That is, the data source produces several (i.e. two) data packets within the 40 ms SDU inter-arrival rate that are grouped into a single SDU. One retransmission is shown for RLC SDUs #1 and #2, both containing 2 VoIP packets. ...”
106. In this example two VoIP packets are grouped into each RLC SDU. The MAC-e PDUs containing RLC SDUs #1 and #2 are the subject of HARQ retransmissions. The MAC-e PDU containing RLC SDU #1 (labelled A) is transmitted in HARQ process number 2 (as shown by the line labelled “HARQ PROCESS ID”). The SDU inter-arrival rate of 40 ms (shown by the vertical arrows) means that the next point at which a MAC-e PDU containing an SDU can be sent is 40ms later, again using HARQ process 2. At this point the MAC-e PDU containing RLC SDU #2 is available for transmission, but there is a retransmission of SDU #1 (in dashed outline and labelled

from the application to MAC every 20 ms. Here packets [sic] #1 (A) and packets [sic] #2 (B) are delivered to MAC layer in their own separate RLC SDUs with a 20 ms time difference, and grouped at MAC-d layer for transmittal in the same single 10 ms air interface TTI.”

112. In this example the VoIP packets arrive separately rather than grouped. Thus VoIP packets are delivered to the RLC, and RLC SDUs each containing one VoIP packet are delivered to the MAC layer, every 20 ms (as shown by the vertical arrows). Since both packets #1 (labelled A) and #2 (labelled B) are small enough to fit into a single MAC-e PDU, they are grouped together at the MAC-d and sent out together in one MAC-e PDU in a single TTI. This first MAC-e PDU is sent on HARQ process 2.

113. The specification continues at [0040]:

“The subsequent new packets #3 and #4 are grouped together in MAC-d for transmittal in the same 10 ms air interface TTI, because the 40 ms virtual TTI prevents this transmission before 40 ms after the start of the previous transmission has elapsed. The packets #1 (A) and #2 (B) are retransmitted (C and D). The packet #1 (A) is transmitted from the source 20 ms earlier than packet #2 (B), but is not allowed to be transmitted before the 40 ms virtual TTI has elapsed after the start of the previous packet first transmission.”

114. Due to the vTTI (shown on the top line of Figure 5), the MAC layer is only allowed to start transmissions of data every 40 ms. Therefore the next point at which a MAC-e PDU can be sent is 40 ms later, again using HARQ process 2. In this example, there is a retransmission of the first MAC-e PDU (containing SDUs #1 and #2 in dashed lines labelled C and D), and so the MAC-e PDU containing the next two packets (#3 and #4, labelled E and F) is delayed until the next TTI, when it is sent on HARQ process 3. At the start of the next vTTI, HARQ process 2 is freed up so the MAC-e PDU containing packets #5 and #6 can be sent on HARQ process 2, followed by a retransmission of the second MAC-e PDU containing packets #3 and #4.

115. The position is therefore similar to that seen in Figure 3: the vTTI defines the cyclic starting point from which a transmission is allowed to occur, but it does not define or determine the length of the gap between the end of one transmission and the start of the next. The precise moment when a MAC-e PDU is in fact sent may be delayed. The actual gaps from the end of one transmission to the start of the next seen in Figure 5 range from 20 to 30 ms, although the vTTI is 40 ms.

116. After the discussion of Figures 5 and 6, the specification states at [0043]:

“Additionally, it is possible to take into account the ‘virtual TTI’ in the application layer to optimize the protocol header overhead, as discussed below. Defining the parameter in the MAC layer advantageously supports the elimination of the dependency on the radio access network, as compared to the case where the transmission interval is defined in the PDP context/RAB parameter. If several radio bearers (RB) are

multiplexed into the same transport channel, it should be possible to separately define a ‘virtual TTI’ for each RB.”


117. After discussing some embodiments which are not material for present purposes, [0048] discloses an embodiment in which two application (e.g. VoIP) packets are inserted into a single UDP/RTP/IP packet in the application layer, in order to reduce the protocol overhead and processing load in the UE and the network.
118. At [0049] the specification points out that having a single speech packet per RLC SDU or UDP/RTP/IP packet has the advantage of greater flexibility, in that this permits the sending of a single speech packet per TTI, for example if transmission was blocked or restricted by higher priority packets, preventing the transmission of large transport blocks that contain several speech packets.
119. The core of the dispute between the parties concerns the next passage in the specification at [0050] to [0052], which provides the basis for the claims as sought to be amended.
120. At [0050] the specification states (I have divided the paragraph into two for clarity and added emphasis):

“In accordance with the present contemplated embodiments, the UE power limitations in bad radio conditions, such as the UE running out of transmission power, are taken into account so that it then becomes possible to send a single speech packet per air interface TTI of 2 ms or 10 ms. Here, the MAC-d would check the RLC buffer of the UE once per virtual TTI, i.e., at the same interval as a normal TTI As a result, packets received during the virtual TTI would be buffered at the RLC level.

In addition, the MAC is permitted to check the RLC buffer more frequently *in certain special cases*, such as when it is not possible to clear the RLC buffer due to power limitations, the transmission of higher priority packets from other RLC buffer[s] (e.g. SRB) or if there are bigger RLC SDUs (e.g., non-compressed headers or real time control protocol (RTCP) packets) that cannot be transmitted within one air interface TTI.”

121. The first half of the paragraph states that the MAC could check the RLC buffer only once per vTTI, in the same way as the RLC buffer is checked only once in a normal TTI. It goes on to say that further packets received during the vTTI have to be buffered in the RLC buffer i.e. until the next vTTI. The skilled reader would understand that the packets stored in the buffer are transmitted in the next vTTI.
122. The second half of this paragraph states that in some circumstances the MAC is permitted to check the RLC buffer more frequently (i.e. more frequently than once per vTTI). The skilled reader would understand this to mean that, in those circumstances, the MAC is able to transmit during more than one TTI within the vTTI.

123. The Defendants emphasise that this is said to be permitted “in certain special cases”. The Defendants contend that this means that this method is an exception to the norm. Conversant does not dispute that the method is described as operating in exceptional circumstances, but points out that this does not alter the method itself. Dr Brydon accepted this point.
124. Paragraph [0050] gives three examples of circumstances when this may apply: when it is not possible to clear the RLC buffer due to (i) power limitations or (ii) the transmission of higher priority packets from other RLC buffers (e.g. SRB or Signalling Radio Bearer packets) and (iii) if there are bigger RLC SDUs that cannot be transmitted within one TTI. It should be noted that these include the scenario where the RLC buffer cannot be cleared due to the sending of other data (e.g. SRB signalling information). The specification therefore expressly contemplates that there may be more than one MAC-d flow and that the vTTI may be implemented across them. This is also recognised in [0052] (set out below) which refers to clearing the “RLC buffers” (plural).
125. This method is illustrated in Figure 7:



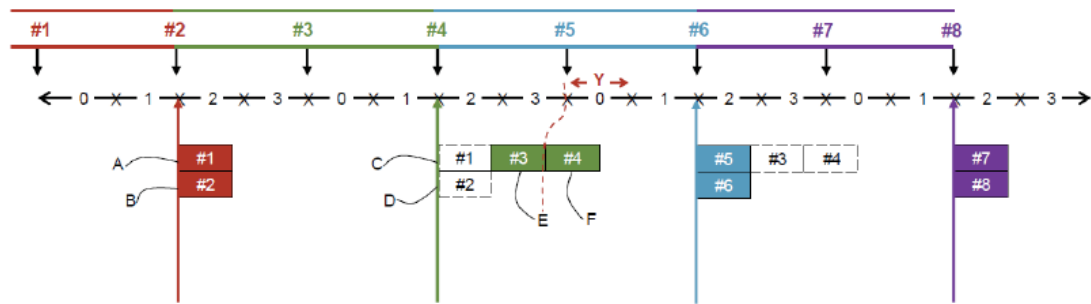
126. The specification describes this at [0051] as follows:

“With reference to FIG. 7, for a 10 ms air interface and a virtual TTI of 40 ms, two VoIP packets are transmitted every 40 ms. Here, packets #1 (A) and packets #2 (B) are grouped together into a single SDU since the subsequent new packet #2 (B) is transmitted within the 40 ms time period. A single transmission for packets #1 (C) and #2 (D) and #3 (E) and #4 (F) is shown. However, it is not possible to transmit subsequent new packets #3 (E) and #4 (F) during a single TTI. As a result, these packets are grouped together and transmitted in separate TTIs.”

127. There is a dispute as to how Figure 7 and [0051] should be interpreted. As the Defendants point out, at first blush, Figure 7 appears to be a slight variant of Figure 5. Moreover, in both Figures, the second line reads “VoIP packets from application (RLC SDUs) every 20 ms”. Still further, both the second sentence of [0039] (describing Figure 5) and the first sentence of [0051] (describing Figure 7) refer to two VoIP packets being transmitted every 40 ms.

128. As Conversant points out, however, the second sentence of [0051] states that “packets #1 (A) and packets #2 (B) are grouped together into a single SDU” i.e. VoIP packets are grouped together in pairs in RLC SDUs at a higher layer, which are then periodically provided to the MAC. This corresponds to the situation in Figure 3 and [0037]. It is in contrast to the position in Figure 5 and [0039], where there is a VoIP packet transmitted every 20 ms and each RLC SDU contains one VoIP packet.
129. As both experts agreed, it follows that there is an inconsistency between the text of [0051] and Figure 7. Counsel for Conversant submitted that the skilled person would give more weight to the text, and conclude that the figure was erroneous. I agree with this, because, as explained below, Figure 7 contains three other obvious errors and one rather more subtle error. In any event, as counsel for Conversant also submitted, it probably does not much matter how the skilled person would resolve the inconsistency, because they would appreciate that the possibilities are those indicated in [0018] (grouping at a higher layer) and [0019] (arriving separately at the MAC layer).
130. VoIP packets #1 and #2 (whether they arrive grouped into one SDU or in separate SDUs) are shown labelled A and B and transmitted in one MAC-e PDU on HARQ process 2. It is common ground that there is an obvious error in the diagram here, because the boxes are shown dashed (indicating a retransmission) when they should be solid (indicating a new transmission).
131. There is then a delay for 40 ms, according to the vTTI of 40 ms, during which time the pair of VoIP packets #3 and #4 arrive. There is a retransmission of #1 and #2 in the next HARQ process (labelled C and D), so the transmission of VoIP packets #3 and #4 cannot begin until the subsequent TTI. VoIP packets #3 and #4 cannot be sent in a single TTI, however. (This could be due to poor conditions, other data or the large size of the RLC SDUs as discussed in [0050].) VoIP packets #3 and #4 are therefore sent in sequential TTIs. It can be seen from Figure 7, although it is not mentioned in the text, that they are retransmitted and therefore delayed by packets #5 and #6 (as in Figure 5).
132. Although not mentioned in [0051], the skilled person would appreciate that, as before, the lines at the bottom indicate the overall delays experienced by the VoIP packets. In the case of #1 and #2, this is 70 ms and 50 ms respectively (as in both Figure 3 and Figure 5). In the case of #3 and #4, it is 80 ms and 50 ms respectively. There is another obvious error in the diagram here, because the third and fourth lines are mislabelled #1 and #2. In addition, the length of the delay shown for #4 is incorrect: it should be 70 ms (10 ms longer than in Figure 5) rather than 50 ms (10 ms shorter).
133. In addition to the obvious errors I have noted, Dr Brydon pointed out that, on close inspection of Figure 7, there is a discrepancy between the length of the space provided for each TTI in the HARQ process line and the size of the boxes depicting the packets. This was not something Prof Marshall had spotted when preparing his first report, and Dr Brydon accepted that Prof Marshall had not been unreasonable or lacking in diligence for not doing so. After Dr Brydon pointed this out, however, Prof Marshall accepted that it appears that all of the boxes in Figure 7 have been drawn slightly too small.

134. Dr Brydon's opinion was that Figure 7 therefore had to be notionally corrected as shown in the following diagram from his first report:



135. In this diagram Dr Brydon has added numbers at the top showing the VoIP packets arriving (this assumes they arrive separately). The dotted red line shows a misalignment of the position of packet #4 when it is transmitted at F. It should be shown as being transmitted on HARQ process 0 in the TTI labelled Y.
136. A number of points arise from this. First, Conversant contends that it would not occur to the skilled person to undertake an analysis of this kind. I disagree with this. Having spotted the inconsistency between [0051] and Figure 7 and the obvious errors in Figure 7, I think the skilled reader would be led to study Figure 7 carefully.
137. Secondly, in his first report, Dr Brydon appeared to be relying upon his correction of Figure 7 when he came to consider the timing of the arrival of packet #5. In cross-examination, however, he explained that this was not the case and that they were independent points. In his diagram Dr Brydon shows the VoIP packets arriving separately, but whether this is correct depends on how the inconsistency between [0051] and Figure 7 discussed above is resolved.
138. Thirdly, Conversant contends that the purpose of Figure 7 is to try to show what happens to packets #3 and #4 as a consequence of the functionality of being able to send these packets over more than one TTI. As noted above, [0051] says nothing about packets #5 and #6 at all. I agree with this, but it does not add anything to the debate over the inconsistency between [0051] and Figure 7.
139. Fourthly, Conversant contends that, whatever the debate about the exact details of [0051] and Figure 7, none of it undermines the clear teaching of [0052]. I will consider that question below.
140. Fifthly, Conversant points out that in Figure 7 one also sees once more that the vTTI cyclically defines the points at which a new transmission may start, but does not indicate the length of the gap between the end of one set of transmissions and the start of the next. The vTTI in Figure 7 is given as 40 ms, but the gaps between the respective transmissions vary from between 10 and 30 ms, due to the effect of retransmissions and sending of VoIP packets #3 and #4 in separate TTIs.
141. The specification goes on at [0052] (again, I have divided the paragraph into two for clarity):

“Examples of the operation of the UE MAC under such conditions are as follows: (i) if the MAC is able to empty the RLC buffer during this air interface TTI, then the MAC will check the RLC buffer at the next predetermined subsequent time interval after the virtual TTI; (ii) if the MAC is not able to empty the buffer, then the MAC will also check the RLC buffer for the next air interface TTI. This permits rapid clearing of the RLC buffers when required, i.e., when large SDUs are utilized. However, the frequency of transmissions during normal operation remains limited.

In certain embodiments, the implementation of the present contemplated embodiment is permitted based on the configuration of the network, e.g., the network is configured to restrict the transmissions only to instances of times that are established by the virtual TTI or the network is configured to permit the above previously described operations.”

142. Again, there is a dispute as to what the first half of this paragraph discloses. Conversant contends that the process described is that, if (i) the MAC is able to empty the RLC buffer in a particular TTI, it will stop transmitting and wait until the next vTTI before transmitting again. If (ii) the MAC is not able to empty the RLC buffer in a particular TTI, then it will check the RLC buffer in the next TTI and transmit the excess data in that TTI. It is implicit that the latter step will be repeated until the RLC buffer is empty. This permits the rapid clearing of the RLC buffers (plural), contemplating the application of the vTTI across several or all of those buffers.
143. The Defendants categorise (i) as “normal operation” and (ii) as exceptional. I have already discussed this distinction in paragraph 123 above. More importantly, the Defendants take issue with what happens in (ii).
144. The case which was put by counsel for the Defendants to Prof Marshall during cross-examination was that option (ii) in [0052] permitted checking in only one more TTI after the first, and no more. Prof Marshall did not accept this. In his view, what the specification was saying was that “If I can empty the RLC buffer in this TTI, I will finish; but if I cannot empty it, I will keep going”. When Dr Brydon was asked about the interpretation which had been put to Prof Marshall, he promptly and unequivocally disavowed it. To the contrary, Dr Brydon accepted that, although the specification only showed one more TTI (i.e. in Figure 7), the skilled reader would understand that there could be (say) two more.
145. Instead, Dr Brydon’s opinion was that the RLC buffer could only be checked for data which was “intended” to have been sent in a vTTI. When questioned on what “intended” meant and how it would be achieved, it turned out to be his way of saying data that had entered in the previous vTTI. When it was put to him that that that would be almost impossible to code for, he made a new suggestion that [0052] was contemplating storing the relevant data in some *different* buffer to the standard RLC buffer and that therefore the UE was not permitted to check the RLC buffer in each TTI (until cleared). As counsel for Conversant submitted, however, not only was that suggestion not put to Prof Marshall, but also it is hard to reconcile with the language of [0052], which simply refers to “the RLC buffer(s)”.

146. Finally, the specification discusses the drawbacks of conventional systems and the benefits of the invention at [0053] and [0054] respectively. The latter paragraph states:

“The present invention advantageously conserves control overhead. Specifically, the downlink E-DCH HARQ Indicator Channel (HICH) (i.e., the HARQ ACK/NAK is sent on E-HICH) overhead can be reduced because ACK/NAKs are needed less frequently. Also, the E-DCH dedicated physical control channel (E-DPCCH) overhead is reduced. In addition, a further savings [sic] in overhead is possible if the uplink dedicated physical control channel (UL DPCCH) gating is introduced. In this case, the DPCCH is not transmitted continuously but only when other UL channels are transmitted. In addition to the savings in system capacity, another advantage of the present invention is that battery power of the UE is conserved, since the UE needs to transmit and receive less often when a virtual TTI is used.”

147. All of these benefits rely on a reduction in the frequency of uplink MAC-e PDU transmissions over the air interface.

The claims

148. As sought to be amended and broken down into integers, claim 1 is to:

- “[A] A method which is executed by a mobile station for autonomous enhanced uplink transmission in which a scheduling grant from a network is not required, comprising:
- [B] determining a virtual transmission time interval for a medium access control entity, which virtual transmission time interval defines a minimum time interval that is allowed between enhanced uplink transmissions;
- [C] checking to determine whether the medium access control entity is transmitting data packets in a current air interface transmission time interval, by checking whether the medium access control entity is able to empty the radio link control buffer in the current air interface transmission time interval;
- [D] and for the case where it is determined that the medium access control entity is not transmitting in the current air interface transmission time interval, transmitting a next data packet only after a period determined by the virtual transmission time interval is determined to have elapsed.”

149. Claim 9 as granted (claim 8 as proposed to be amended) is to:

“The method of claim 1, wherein determining the virtual transmission time interval comprises receiving from a network element the virtual transmission time interval.”

150. Claims 12 and 20 as granted (claims 11 and 18 as proposed to be amended) are parallel computer program claims. Claims 23 and 31 as granted (claims 21 and 28 as proposed to be amended) are parallel mobile station claims. Granted claims 4, 15 and 26 are to be deleted.

Construction

151. There is a series of disputes as to the interpretation of claim 1. There is no dispute that the claim is to be interpreted purposively in the light of the specification. As counsel for Conversant submitted, the issues of construction should be considered separately from the question of whether, once construed, the claims read on to Uplink DRX, rather than rolling up the two sets of issues together, as counsel for the Defendants did. Nevertheless, there is one issue of construction, that relates to the Defendants’ case of obviousness over Samsung, which it is more convenient to discuss in that context.

[A] A method ... for autonomous enhanced uplink transmission

152. The Defendants contend that this means that the method *selectively* limits the transmission rate for autonomous transmissions, but *without affecting* the UE’s ability to make non-autonomous i.e. scheduled transmissions. Conversant contends that the claim is only concerned with what happens to autonomous transmissions and does not address what happens to scheduled transmissions.
153. In my judgment Conversant is correct. As counsel for Conversant submitted, the language of the claim neither requires a selective effect on the autonomous transmissions nor specifies that the scheduled transmissions are unaffected.
154. Counsel for the Defendants argued that it would be contrary to the purpose of the invention to slow down scheduled transmissions, but I do not accept this. There is no dispute that the problems which the Patent identifies are problems with autonomous transmissions or that the invention aims to address those problems. Nor is there any dispute that the invention is particularly concerned with reducing the margin which needs to be reserved for autonomous transmissions. Nor is there any dispute that the invention involves reducing the frequency at which MAC-e PDUs containing data being transmitted autonomously are delivered to the PHY. It does not follow that scheduled transmissions must be left unaffected. On the contrary, Dr Brydon accepted that overhead benefits could be obtained by grouping packets together and thus reducing signalling and headers. This applies to both autonomous and scheduled transmissions.
155. Furthermore, as discussed above, I consider the skilled reader would understand from the specification that the method disclosed in [0015], [0032] and Figure 2 involved slowing down scheduled transmissions as well as autonomous transmissions. Even if that is wrong, the skilled reader would understand that the method disclosed in [0016] and [0033] involved slowing down some scheduled transmissions as well as autonomous transmissions.

[B] a virtual transmission time interval for a medium access control entity

156. The Defendants contend that this means that the vTTI is to be implemented in the MAC-d and to affect MAC-d flows, relying in particular upon what the specification says at [0019]. Conversant disputes this.
157. In my judgment Conversant is correct. As Conversant points out, the wording of the claim is entirely general, referring to “a” MAC entity. It is not restricted to MAC-d. Furthermore, in so far as the Defendants’ interpretation is based on the meaning which they say the skilled person would attribute to the expression “virtual transmission time interval”, it would render the limitation “for a medium access control entity” otiose. Still further, [0019] is describing a specific embodiment. Moreover, as discussed above, it describes what is possible, not what is mandatory.
158. The Defendants also contend that the method must be service specific. Conversant disputes this.
159. In my judgment Conversant is again correct. There are simply no words in the claim which limit it to a service specific solution. Furthermore, such a limitation would be contrary to the teaching of the specification, where non-service specific examples are given in [0015]/[0032] and [0016]/[0033]. Even in the case of the MAC-d implementation in [0019], as discussed above, a MAC-d flow may contain more than one service. Still further, Dr Brydon accepted that the skilled person would understand that the Patent was teaching that the vTTI could be applied to more than one service or more than one MAC-d flow.

[B] which virtual transmission time interval defines a minimum time interval that is allowed between enhanced uplink transmissions

160. The Defendants contend that the vTTI defines a fixed gap between transmissions. Conversant disputes this. It contends that the vTTI defines the permitted starting points for transmissions. In my judgment the language of the claim can be read in either way, but it is clear from the specification, and in particular Figures 3, 5 and 7 as discussed above, that Conversant is correct.

[C] data packets

161. Integer [C] refers to the MAC transmitting “data packets”. The Defendants contend that these must be VoIP packets. As Conversant points out, however, there is no such limitation in the claim. The expression “data packets” is entirely general. Moreover, as Prof Marshall noted, the term is used in a variety of contexts in the specification.

[C] checking to determine whether the medium access control entity is transmitting data packets in a current air interface transmission time interval, by checking whether the medium access control entity is able to empty the radio link control buffer in the current air interface transmission time interval

162. Conversant contends that integer C needs to be read in the light of integer D. The check in integer C is binary in nature i.e. it will return a positive or negative result. If the result is negative, then integer D says that the UE must wait until the end of the

vTTI to transmit again. Thus far I do not think that there is any dispute, but in any event I agree with this.

163. The dispute concerns a transitional TTI in which the UE moves from one that is transmitting at the beginning of the TTI to one that is not transmitting at the end of the TTI. Does that mean that the UE is transmitting in the current TTI or does it mean that the UE is not transmitting in the current TTI? The answer depends on when the check is carried out: if the check is carried out at the beginning of the TTI, the MAC is transmitting; but if the check is carried out at the end of the TTI, the MAC is not transmitting.
164. Conversant contends that the skilled person would understand that the purpose of integer C is to identify the relevant TTI so as to apply integer D, and that this is done by checking whether the MAC is able to empty the RLC buffer in the current TTI. Conversant further contends that the skilled person would understand from this that the patentee intended that, where the MAC is no longer transmitting at the end of a particular TTI, the check would yield a negative answer. In other words, if the MAC is able to empty the RLC buffer by the end of the TTI, then that indicates that the MAC is no longer transmitting at that point.
165. As I understand it, the Defendants dispute this for two reasons. First, they contend that the language of the claim indicates that the check is whether the MAC is in a transmitting state at any point during the TTI. It is immaterial whether the MAC finishes transmitting at the end of the TTI.
166. Secondly, the Defendants contend that it not possible to determine whether the MAC is transmitting by considering whether it is able to empty the RLC buffer. They say that whether or not the RLC buffer is able to be cleared is a different test to whether it is transmitting or not.
167. In my judgment Conversant is correct that, as a matter of construction of the claim as proposed to be amended, integer C involves checking whether the MAC is transmitting at the end of the current TTI. That is consistent with the check being carried out by checking that the MAC is able to empty the RLC buffer. If it were otherwise, then the Defendants would be correct that the check as to whether the MAC is able to empty the RLC buffer would not determine whether or not the MAC is transmitting. As will appear, however, it is also necessary to consider separately an issue as to added matter which arises (see paragraphs 219-227 below).

Essentiality

Uplink DTX

168. Uplink Discontinuous Transmission (Uplink DTX or UL DTX) was introduced in Release 7 of UMTS. It is a mechanism of discontinuous transmission on the legacy DCH control channel, the Dedicated Physical Control Channel (DPCCH). In normal circumstances control data is sent on the DPCCH every TTI. Where, however, there are no transmissions made on the DCH data channel (the Dedicated Physical Data Channel or DPDCH) for a period, then the Uplink DTX kicks in and no transmissions are generally made on the DPCCH either. It is necessary, however, to send periodic

bursts to ensure resources are not torn down by the network, and other reasons such as power control.

169. Once entered into, Uplink DTX limits when bursts can be sent on the DPCCH. These are restricted to particular fixed periods with fixed start and end points. Bursts are also sent on the DPCCH when there are transmissions of data on the E-DCH data channel, the E-DPDCH (as well as corresponding bursts on the enhanced dedicated physical control channel – the E-DPCCH). These DPCCH bursts are not, however, limited by the Uplink DTX and may be sent at any time.

Uplink DRX

170. Uplink DRX or UL DRX is a mechanism for discontinuous transmission on the E-DCH.
171. When Uplink DRX is active, the MAC-e is restricted in when E-TFC selection may be performed and therefore when a MAC-e PDU can be generated. There are two key aspects of the system, the “MAC Inactivity Threshold” and the “MAC DTX cycle”.
172. Section 11.8.1.4 of TS 25.321 (v9.5.0 December 2010) states:

“In FDD, for each Activated Uplink Frequency, in case the DTX feature is configured by higher layers and no E-DCH transmission is performed in this TTI on that Activated Uplink Frequency:

- if *MAC Inactivity Threshold* > 1 and no E-DCH transmission has been performed for *MAC Inactivity Threshold - 1* previous TTIs or,
- if *MAC Inactivity Threshold* = 1:
 - E-TFC selection shall only be performed for the TTIs where the following conditions are fulfilled:
 - For 2ms TTI: $[5 * CFN + \text{subframe number} - UE DTX DRX Offset] \bmod MAC DTX Cycle = 0$;
 - For 10ms TTI: $[5 * CFN - UE DTX DRX Offset] \bmod MAC DTX Cycle = 0$.

173. The process operates as follows. The MAC-e monitors E-DCH transmissions. If in a certain TTI an E-DCH transmission is performed, then the MAC can proceed to the next TTI where it can perform a transmission, if needed, without regards to the MAC DTX cycle.
174. If, however, in that TTI there is no E-DCH transmission, then the indented pseudocode is engaged. First the MAC refers to the “MAC Inactivity Threshold”. This is a parameter which can be set to 1, 2, 4, 8, 16, 32, 64, 128, 256, 512 or infinity units of E-DCH TTIs.

175. Conversant's case is based on the scenario where MAC Inactivity Threshold = 1. The Defendants suggest that this is an artificial and contrived case because, although the Standard requires phone manufacturers to cater for the possibility that MAC Inactivity Threshold is 1, it is very unlikely that MAC Inactivity Threshold would be set by the network to 1 (according to the Defendants, it is typically set to 64). No non-infringement case was advanced by the Defendants on this basis, however. On the contrary, the Defendants have admitted that, if the Patent is essential and valid, then they have infringed it. As Conversant points out, it follows that what happens if MAC Inactivity Threshold is greater than 1 is irrelevant.
176. In cases where the MAC Inactivity Threshold is set to 1, then the next section of pseudocode is engaged, which states that E-TFC selection shall only be performed for the TTIs where the specified condition is fulfilled. This involves the parameter MAC DTX cycle.
177. The MAC DTX cycle defines on a cyclical basis the next TTI in which E-TFC selection can be performed (and therefore the next MAC-PDU transmitted) for E-DCH transmissions.
178. The formula functions as follows:
- i) The CFN (Connection Frame Number) and subframe numbers are numbers which are assigned sequentially to identify a particular frame and subframe. The UE DTX DRX Offset is a parameter which is set by the network in effect to space when the respective MAC DTX cycles of different UEs allow transmissions to occur.
 - ii) Thus for a 2ms TTI, E-DCH transmissions are only allowed when $(5 * \text{current Connection Frame Number} - \text{UE DTX DRX Offset} + \text{current subframe})$ is an integer multiple of MAC DTX cycle. This is because the **mod** operator finds the remainder after division of one number by another.
179. The effect of this is that, if an E-DCH transmission occurred in the previous TTI, then the UE can transmit in the current TTI; but if there was no E-DCH transmission in the previous TTI, then the MAC must wait until the arrival of the next TTI in which the MAC DTX cycle permits E-TFC selection before any new E-DCH transmission can be performed. If a new transmission is made during a TTI that is permitted by the MAC DTX cycle, the restriction is lifted and a further transmission can be made in the next TTI.

Buffer Occupancy

180. Conversant's pleaded case as to the way in which the check at integer C of claim 1 is carried out refers to a defined parameter called Buffer Occupancy (BO). There was some confusion during the cross-examination of the experts between this specific parameter and the concept of buffer occupancy more generally, but the pleaded case is clear.
181. Conversant's pleaded case is that the check involves the MAC using the BO parameter to obtain an indication of the amount of data available for transmission in the RLC buffer. Prof Marshall proceeded upon the same basis in his reports.

182. In cross-examination it was put to Prof Marshall, and he accepted, that the Standard does not mandate use of the BO parameter to determine whether an E-DCH transmission has been performed. He acknowledged that it would be possible to check buffer occupancy in another way. For his part, Dr Brydon accepted that the E-TFC selection process “will take account of the amount of data in the buffers in some way”, but maintained that this did not necessarily involve use of the BO parameter. He opined that it was not appropriate for that purpose, and agreed with Prof Marshall that it could be done in other ways.
183. This evidence gave rise to a major dispute in closing submissions. Counsel for the Defendants submitted that Conversant was advancing a new case which was not open to it, and which in any event was unsubstantiated because no specific alternative to the BO had been identified. Counsel for Conversant disputed this. He accepted that the evidence established that use of the BO parameter was not *mandatory*, but submitted that it did not establish it was not *possible*. He also pointed out, correctly, there was no evidence that anything other than the BO parameter was in fact used: nothing else was put to Prof Marshall and nothing else was identified by Dr Brydon. In those circumstances, he submitted that it was probable that BO was used. He also submitted that it made no difference if the same information was conveyed in a different way, because the key point, which Conversant had maintained all along and Dr Brydon had accepted, was that the Standard required the MAC to check the data available in the buffer and then indicate whether there was a need for a new transmission.
184. I accept the submissions of counsel for Conversant on this point. On the evidence, it appears more likely than not that the BO parameter is used. In any event, I agree that it does not matter if the same information is conveyed in a different way, because it makes no difference to Conversant’s essentiality case. In this regard, I note that Conversant made this very point in paragraphs 153-155 and 161-162 of its opening skeleton argument, but no objection was taken by the Defendants at that stage that Conversant was advancing a new and unpleaded case.

Retransmissions in Uplink DRX

185. There is also a more minor dispute about the role of retransmissions in Uplink DRX. Dr Brydon explained his view of the operation of the Uplink DRX cycle through four worked examples A-D. Prof Marshall largely agreed with Dr Brydon’s explanation, but there was some disagreement as to Dr Brydon’s Example C, relating to one aspect of how retransmissions are handled.
186. The issue is whether a retransmission in an allowed TTI has the same effect as a new transmission would have (i.e. to end the period of restriction) or whether a retransmission does not lift the restriction such that transmissions are effectively blocked again for another MAC DTX cycle.
187. Prof Marshall gave evidence, which was not challenged in cross examination, that a retransmission in an “allowed” TTI would lift the restriction in the same way as a new transmission. As he pointed out, that is consistent with the description of how the MAC DTX cycle works in the Standard at section 11.1.1 of TS 25.308:

“Whenever the UE transmits one E-DCH TTI the UE can use subsequent TTIs for E-DCH transmission as long as its transmission is continued (re-started) within UE_Inactivity_Threshold TTIs, and in addition the pending retransmissions can be done in the corresponding HARQ processes (retransmission times are not impacted by the inactivity threshold).”

188. Dr Brydon acknowledged that the pseudocode in the MAC Standard (which he agreed was in any case imprecise) did not make any explicit provision for how the restriction imposed by the MAC DTX cycle would be lifted. He also recognised that, if the Standard operated in accordance with his understanding, it would result in HARQ retransmissions blocking any new transmissions for several MAC DTX cycle periods, which would be entirely contrary to the principle behind using several HARQ processes of avoiding a “stop and wait” error correction system. He also acknowledged that it would be contrary to the way in which the MAC DTX cycle is described in TS 25.308. Most importantly, Dr Brydon acknowledged that the distinction probably had no real significance with respect to essentiality.
189. Counsel for Conversant submitted that Prof Marshall’s view was supported by the following considerations. First, it is common ground that the pseudocode does not expressly state when the period of restriction comes to an end, but it is implicit that it is intended that periods of restriction do end. Prof Marshall’s interpretation is the only way to read the system as doing what is described generally in the standard in 11.1.1 of TS 25.308. Those passages make it clear that E-DCH transmissions (i.e. the category comprising both new and retransmissions) are to be provided with a pattern of start times. If a re-transmission takes place in a TTI provided as one of the start TTIs, E-DCH transmissions do start at that point. Secondly, Prof Marshall’s interpretation would not have the unnecessary (and highly undesirable) result of blocking of the system completely, possibly for several complete DTX cycles – something that HARQ is specifically intended not to do.
190. I accept these submissions. Accordingly, I accept Prof Marshall’s evidence on this issue.

Conversant’s case in outline

191. *Integer A.* The method is Uplink DRX, which is executed in a mobile device. E-DCH transmission of non-scheduled data is an autonomous transmission for which a scheduling grant from the network is not required.
192. *Integer B.* The ν TTI is the MAC DTX cycle. This is implemented in the MAC. Once triggered, the MAC DTX cycle prevents the making of any E-DCH transmissions (but not retransmissions) until they are allowed in accordance with the MAC DTX cycle. It provides, therefore, a cyclic pattern of start points which defines the minimum time interval that is allowed between E-DCH transmissions, in accordance with the claim.
193. *Integers C and D.* Considering a situation where the MAC-e entity is allowed to transmit in a particular TTI because, for example, it was transmitting in the previous TTI:

- i) If the reported buffer occupancy is not zero, then the E-TFC selection entity will indicate to the HARQ entity that there is a need for a new transmission. As a result, data will be pulled down from the RLC buffer, a MAC-e PDU will be constructed, and sent on to PHY. The MAC therefore determines that it will transmit a MAC-e PDU in this TTI by reference to the buffer occupancy.
- ii) If, however, the buffer occupancy is zero, then the E-TFC selection entity will not indicate the need for a new transmission, so no data will be pulled down from the RLC buffer, and no MAC-e PDU be constructed or sent on to the PHY. The MAC therefore determines that it will not transmit a MAC-e PDU in this TTI by reference to the buffer occupancy. Because of this, no E-DCH transmission is performed in that TTI, and the restriction under section 11.8.1.4 will kick in. The MAC will therefore be unable to transmit a new MAC-e PDU on the E-DCH, until after the period determined by the MAC DTX cycle has elapsed.

194. Conversant accepts that, even on its construction of the claim, there is no literal infringement because the MAC is checking the status of the RLC buffer by reference to whether the RLC buffer was able to be emptied in the *previous*, rather than the *current*, TTI. Conversant contends that there is infringement on the basis of the doctrine of equivalence expounded by the Supreme Court in *Actavis UK Ltd v Eli Lilly & Co* [2017] UKSC 48, [2017] RPC 21.

Assessment

195. *Integer A.* It is common ground that E-DCH transmissions include both scheduled and non-scheduled data. The Defendants contend that this means that Uplink DRX does not comply with integer A because, as a matter of construction, integer A requires that the method applies selectively to autonomous transmissions and not to scheduled transmissions. Since I have not accepted that construction of the claim, this integer is satisfied.
196. *Integer B.* The Defendants contend that this integer is not satisfied for two reasons, both of which are matters of construction. First because the claim requires that the vTTI be implemented in the MAC-d. Secondly, because the vTTI defines a minimum gap between transmissions. Since I have not accepted either of these points, this integer is satisfied.
197. *Integer C.* There are a number of issues concerning this integer.
198. First, the Defendants contend that, as a matter of construction, the claim is limited to VoIP packets. I have rejected this interpretation above.
199. Secondly, the Defendants contend that the integer requires a check as to whether there are any transmissions at all, including retransmissions. Conversant contends that this is wrong, and that the transmissions that are relevant are those of the category of data to which the method is applied. I accept Conversant's submission, although I shall have to return to the question of retransmissions below.
200. Thirdly, there is the point on buffer occupancy which I have discussed above.

201. Fourthly, there is the point on equivalents noted above. I found the Defendants' position on this unsatisfactory. Conversant's opening skeleton argument made it clear that Conversant's case was one of equivalence with respect to the timing of the check. It relied upon the evidence given by Prof Marshall in his reports with respect to the three *Actavis* questions. It appeared to be implicit that Conversant accepted that there was no literal infringement, and counsel for Conversant confirmed this in his opening submissions. Counsel for the Defendants did not cross-examine Prof Marshall on the *Actavis* questions, and said nothing at about equivalence in his written closing submissions. When I asked counsel for the Defendants if I was right in understanding that the Defendants took no independent point on equivalence, he initially appeared to confirm this, but subsequently produced supplemental written submissions contending that there was no infringement by equivalence. As with many of the Defendants' other submissions, I found these hard to understand, but in so far as I do understand them, they raise points which I have already considered under other headings, particularly with regard to buffer occupancy. The Defendants have no coherent answer to Conversant's case that Uplink DRX is equivalent to integer C in that the check is carried out with respect to the *previous* TTI rather the *current* TTI but this achieves substantially the same result in substantially the same way as the invention, that this would be obvious to the skilled person and that the skilled person would not think that strict compliance was required. I therefore accept Conversant's case on this point.
202. Fifthly, the Defendants raise a point concerning retransmissions. As discussed above, the Patent says that retransmissions are unaffected: they are "performed normally". Conversant says that Uplink DRX does not affect retransmissions either. The Defendants argue that Uplink DRX is different, because in Uplink DRX retransmission has the effect of allowing transmission in the next TTI. As Dr Brydon accepted by reference to an example (X/1) which was put to him in cross-examination, however, in the ordinary case retransmission has no such effect. As counsel for Conversant submitted, this would not prevent infringement in any event if the claim is construed in the manner that I have accepted.
203. Sixthly, the Defendants raise points concerning two categories of control-only transmissions, namely SRB transmissions and Signalling Information (SI) transmissions. The SRB is no different to other data held in the RLC buffer: it is just a specific radio bearer, but carrying signalling from higher layers. Accordingly, its packets are within Conversant's infringement case anyway and they are treated just like any other radio bearer. As for the SI transmissions, even though they themselves do not comprise data that is ever held in the RLC buffers, they are only sent when there is data in the RLC buffer, as Prof Marshall explained: see section 11.8.1.6 of TS 25.321. In any event, as with retransmissions, restricting additional transmissions does not matter to the issue of infringement as I have construed the claim.
204. *Integer D*. The Defendants contend that this integer is not satisfied either. As I understand it, however, the Defendants do not rely upon any additional arguments in this respect to those which I have already considered above.
205. *Conclusion*. For the reasons given above, I conclude that the Patent is essential to Uplink DRX in UMTS, in that Uplink DRX falls within the scope of protection of claim 1 in the case where the MAC Inactivity Threshold is 1, and the Defendants have therefore infringed at least claim 1 if it is valid.

Added matter

The law

206. There is no dispute as to the applicable legal principles. Added matter requires a comparison of the disclosures of the application as filed and the patent, viewed through the eyes of the skilled addressee. The comparison is strict in the sense that subject matter will be added unless such matter is clearly and unambiguously disclosed in the application, but that may be done explicitly or implicitly. For this purpose, the claims form a source of disclosure as well as the specification. As the Defendants emphasise, one type of added matter is intermediate generalisation, which occurs when a feature is taken from a specific embodiment, stripped of its context and then introduced into the claim when it would not be apparent to the skilled person that it was generally applicable to the invention: see *Nokia Corp v ICom GmbH & Co KG* [2012] EWCA Civ 567, [2013] RPC 5 at [56]-[60] (Kitchin LJ). As Conversant emphasises, it is wrong to equate what a claim discloses with what it covers. The law does not prohibit the addition of claim features which state in more general terms that which is described in the specification. What the law prohibits is the disclosure of new information about the invention: see *AP Racing Ltd v Alcon Components Ltd* [2014] EWCA Civ 40, [2014] RPC 27 at [28]-[40] (Floyd LJ).

Assessment

207. Save for the claims, there is little difference between the Application and the Patent. The main difference is that the Patent omits the following text from the Application at page 18 lines 14-21:

“For example, it is expressly intended that all combinations of those elements and/or method steps which perform substantially the same function in substantially the same way to achieve the same results are within the scope of the invention. Moreover, it should be recognized that structures and/or elements and/or method steps shown and/or described in connection with any disclosed form or embodiment of the invention may be incorporated in any other disclosed or described or suggested form or embodiment as a general matter of design choice.”

208. Dr Brydon accepted that this is a disclosure that the delaying mechanism depicted in one embodiment is intended to be taken as disclosed and applicable as a delaying mechanism in all embodiments.

209. The Defendants’ pleaded case advances five allegations of added matter which I will consider in turn.

210. *Allegation 1.* Claim 1 of the Application required “checking to determine whether all PDUs are transmitting data packets autonomously”. It is alleged that, if the claims of the Patent do not disclose the same checking, they add matter. Claim 1 of the Application also required “decelerating packet transmission rates”. It is alleged that the claims of the Patent do not disclose the same and add matter.

211. The two points advanced in this plea are of similar nature. Each is an argument that a limitation in the claims of the Application must also be a limitation of the Patent, and that if there is no such limitation, matter is added.
212. Conversant contends that this is wrong for three reasons. First, a granted claim may be broader than or different from the claims filed in the application. The objection is properly concerned with matter that is added. This is not an argument of added matter at all. Secondly, the argument appears to amount to saying that a feature of the Figure 2 embodiment must be a limitation of the claim, but the Figure 3 and Figure 5 embodiments operate differently. Thirdly, so far as deceleration is concerned, the features of the claim do achieve deceleration for those UEs that stop transmitting. The deceleration depends upon whether the MAC can clear the RLC buffer, but that reflects the embodiment disclosed at [0050]-[0052].
213. In my judgment Conversant is correct. The Patent does not disclose any new information about the invention in this respect, and therefore there is no added matter.
214. *Allegation 2.* The second allegation relies upon claims 12 and 13 of the Application. The former introduced a requirement for a new MAC parameter in MAC-d. The latter provided that the parameter is a vTTI. The complaint is that there is no disclosure of a vTTI generally for MAC. In particular it is said it must be implemented in MAC-d, and if the claims extend to a method in which the vTTI is implemented other than in MAC-d, in particular in the MAC-e, matter is added.
215. Again, Conversant has three answers to this. First, Dr Brydon accepted that the skilled person would not understand the place of implementation to be important. Indeed, implementation in the MAC-e was the norm. Secondly, the disclosure in the Application of the idea of a vTTI is not limited to implementation in MAC-d. For example, the mobile station claims in the Application comprise claims which expressly incorporate the vTTI without any limitation as to where it is implemented. Indeed, as noted above, Dr Brydon accepted that the Application disclosed that delaying mechanism could be used in any embodiment. Third, this argument also confuses the scope of a claim and its disclosure. Even if the Defendants' points were otherwise correct, the Patent still does not teach the skilled addressee anything new.
216. In my judgment Conversant is correct. The Patent does not disclose any new information about the invention in this respect, and therefore there is no added matter.
217. *Allegation 3.* The third allegation is that, if the claims require a checking step for other than the first "and next" TTI, then they add matter since the disclosure of the Application would be understood as limited to checking in the first "and next" TTI (i.e. the first and second TTIs only).
218. Conversant contends that this objection is based upon a misunderstanding of the disclosure. As discussed above, the disclosure is not of a limited ability to check only in the second TTI. It is plain that the Application is contemplating a process of checking of the RLC buffer for the purposes of emptying the RLC buffer. Conversant also contends that this objection again muddles claim coverage with disclosure. The claims in the Patent do not disclose that the check should be made otherwise than in the second TTI, although they cover that possibility.

219. *Allegation 4.* This allegation is directed first to the introduction into integer C of the check as to whether the MAC is transmitting in a current TTI and into integer D of the determination that it is not, and secondly to the amendment to integer C. Importantly, the second part of the allegation is raised by the Comptroller in opposition to the amendment application. (The Comptroller is, of course, not concerned with the amendments to the claims made during prosecution.)
220. It is convenient to focus the discussion on integer C. As proposed to be amended, integer C requires “checking to determine whether the medium access control entity is transmitting data packets in a current air interface transmission time interval, by checking whether the medium access control entity is able to empty the radio link control buffer in the current air interface transmission time interval”. The disclosure in the Application which is relied upon as providing the basis for this integer, including the amendment, is the passage corresponding to [0050]-[0052], and in particular [0052].
221. The Defendants and the Comptroller contend that what the Application discloses is checking of the RLC buffer (after the vTTI or the next TTI respectively) to determine if the MAC is either able or not able to empty the RLC buffer; but not checking to determine whether the MAC is transmitting data packets in the current TTI by checking whether the MAC is able to empty the RLC buffer in the current TTI.
222. The Defendants add that the real vice was the amendment of the claim during prosecution to require (and hence disclose) a check to determine whether the MAC is transmitting data packets in a current TTI which was not disclosed in the Application, and that this is not cured by the amendment. (The Defendants also argue that the situation is compounded by Conversant’s equivalence argument, which dilutes (the Defendants go so far as to say as to say “erases”) the new limitation in integer C. In my judgment, however, this argument is misconceived, since it is an argument about the scope of protection, and not about disclosure. Infringement by equivalence cannot give rise to an added matter objection.)
223. Conversant contends that, whatever might be said about the granted claim, this objection ignores the totality of the disclosure of integers C and D of the claim as sought to be amended. All that is disclosed is a specific check to determine whether the consequence of integer D is to apply. The check is same as the check that is disclosed in the Application – the ability to empty the buffer. The purpose of the check is the same, and so is the consequence. Accordingly, Conversant contends, the skilled person learns nothing new about the invention from the amended claim compared to the Application.
224. In considering this issue, the starting point is that there is clearly no explicit disclosure in the Application of a check to determine whether the MAC is transmitting in the current TTI. It follows that there is no explicit disclosure of a check to determine whether the MAC is transmitting in the current TTI by checking whether the MAC is able to empty the RLC buffer in the current TTI. But there is an explicit disclosure of checking whether the MAC is able to empty the RLC buffer during the current TTI. It follows, as I see it, that the key question is what the skilled reader of the Application would understand the purpose and effect of that check to be. In my view the skilled reader would understand that the purpose and effect is as described in paragraphs 142-

145 above. In short, the check is concerned with whether the MAC is able to empty the RLC buffer or not, and the consequences of that.


225. In my judgment it follows that the claim, both as granted and as amended, does disclose new information about the invention. Whereas the passage in the Application corresponding to [0052] discloses checking whether the MAC is able to empty the RLC buffer during the current TTI, and then either waiting to transmit until the next vTTI or transmitting in the next TTI, it does not disclose a check to determine whether or not the MAC is transmitting in the current TTI. This is not altered by the fact that, by the amendment, the latter check is to be carried out by checking whether the MAC is able to empty the RLC. This is not a question merely of what the claim covers, but of what the claim discloses.
226. I therefore conclude that claim 1 as granted is invalid for added matter and that the amendment does not cure the invalidity. Conversant does not contend that, in that event, any other claim is independently valid whether as granted or as proposed to be amended.
227. *Allegation 5.* The final objection is a complaint that the amended claims introduce the concept of checking whether MAC can empty the buffer removed from the context of the “certain special cases” in which it is disclosed in the Application. Conversant argues that this is a bad point. The method is no different in the Application and the amended Patent. The way the UE is set up and configured is precisely the same as in the Application. In my judgment Conversant is correct about this.

Obviousness

228. The Defendants contend that the claimed inventions are obvious in the light of two items of prior art: (i) document R1-041087 submitted by Samsung to the 3GPPP TSG RAN Working Group 1 meeting in Seoul, Korea on 20-24 September 2004 (“Samsung”); and (ii) US Patent Application 2003/00866381 entitled “System for Improved Power Savings During Full DTX Mode of Operation in the Downlink” filed on 18 October 2002 and published on 8 May 2003 (“Terry”). There is no dispute as to the applicable legal principles.

Samsung

229. *The disclosure.* Samsung describes a method for dealing with autonomous transmissions (sent on the minimum set, as was the proposal then), which it describes as those for which no scheduling command is required from the Node B. Samsung identifies a problem with such transmissions, which is that the Node B has to reserve resources to account for the noise rise contributed by autonomous transmissions. This issue is illustrated by Figure 2 of Samsung, which shows a large portion of the overall noise rise contributed to by autonomous transmissions (shown in pink).



230. Samsung states in paragraph 2.2 that the minimum set for autonomous transmission should be minimised, and that the size of transport blocks belonging to the E-DCH minimum set should be as small as possible. Samsung goes on in paragraph 2.2 to state that a method is needed “to lower the effective data rate for” autonomous transmissions. The document then describes three previous schemes. Each of these schemes is said to have severe drawbacks.
231. The method proposed by Samsung is set out in section 2.3. It is to introduce a form of what it describes as Time Division Multiplexing (TDM), by which autonomous transmissions for each UE are restricted to a sub-set of allowed sub-frames (there being 5 sub-frames, each of 2 ms, within each frame of 10 ms). With a 2 ms TTI, each sub-frame corresponds in time to 1 TTI. The approach is illustrated in Figure 1 of Samsung:

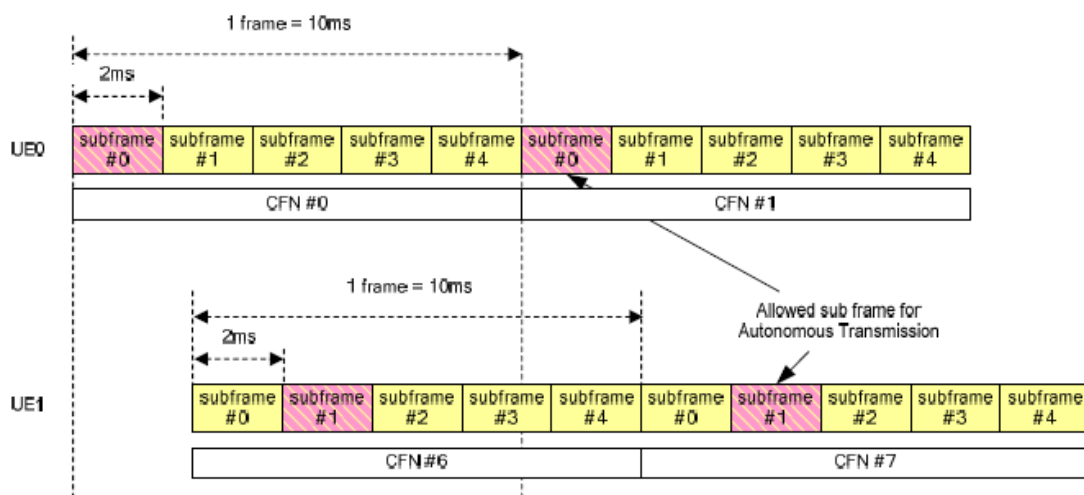


Figure 1: Autonomous transmission with allowed TTI restriction

232. As described by Prof Marshall, this approach effectively applies a “mask” which overlays each UE’s allocation of transmission opportunities for autonomous transmission within the 10ms time frame. He explained that it is a “static” method in that a UE that is restricted to the use of sub-frame 0 cannot under Samsung be permitted to use any of sub-frames 1 – 4 to send more data. Similarly, the fixed sub-frame structure does not allow a reduction in transmission opportunities beyond 1 in 5

sub-frames. Finally, since there are 5 sub-frames, there have to be 5 HARQ processes associated with them (which was not the case in the Standard at the filing date).

233. There is a discussion in Samsung of offsetting the allowed sub-frame for different UEs, to spread out the allowed time for autonomous transmissions for different UEs in the cell. As Prof Marshall noted, to get the full benefit of such an approach, a degree of synchronisation would be required, which was not the case in UMTS uplink and which would add a further technical burden.
234. Samsung then sets out, in section 2.4, the results of a simulation to determine the effect of its method on the noise rise problem. As shown in figures 2 and 3, by reducing the capacity of each UE to transmit autonomous data, the reserved margin due to autonomous transmissions is reduced. As Dr Brydon explained, this allows a larger proportion of the cell's capacity to be used for scheduled transmissions – an effect that is more pronounced as the number of users in the cell is increased.
235. *Obviousness.* Apart from a statement that the use of 1 in 5 sub-frames might be changed to 1 in 6 or 1 in 10, the totality of Dr Brydon's evidence as to obviousness over Samsung in his first report was the following paragraph:
- “Samsung addresses the same problem as the Patent and proposes a method for reducing the rate at which autonomous transmissions are made, in order to reduce their impact on uplink resources, which uses the same principle as the Patent. The Samsung approach is to undertake autonomous transmissions in a fixed cycle of sub-frames. It therefore imposes a ‘minimum time interval’ specifically between autonomous transmissions, which is the same concept as disclosed in the Patent.”
236. In his second report, Dr Brydon responded to some of Prof Marshall's comments on how Samsung operated, but did not add to what he had said in the paragraph quoted above.
237. Counsel for Conversant did not cross-examine Dr Brydon on this evidence. As counsel submitted, there is no attempt by Dr Brydon to identify any modification to Samsung which it would be obvious for the skilled person to make. In particular, Dr Brydon says nothing about the ability of the UE to keep transmitting if the buffer is not cleared.
238. When cross-examined Prof Marshall was clear that all Samsung did was to provide for when the transmission opportunities were; it did not provide for continued transmissions. It was not suggested to Prof Marshall that the skilled reader would modify Samsung to introduce the feature of allowing UEs that were unable to empty their buffer to continue to transmit in consecutive TTIs. Moreover, he was not challenged on his written evidence that this is not contemplated by Samsung.
239. The case argued by counsel for the Defendants in closing submissions was that the only difference between Samsung and claim 1 was that Samsung did not clearly and unmistakably disclose implementation in a MAC that carried out E-TFC selection,

but that that was obvious (as Prof Marshall accepted). Everything else in claim 1, counsel argued, was disclosed.

240. This argument gives rise to an issue which is one of construction, but which it is convenient to deal with here. The issue is what happens if it is determined in integer C that the MAC *is* transmitting in the current TTI. On the face of it, integer D of the claim is silent as to what should happen in that event. Conversant contends that it is implicit that the MAC should carry on transmitting. The Defendants dispute this.
241. Counsel for Conversant complained that this argument had only been raised by the Defendants for the first time in their written closing submissions. Counsel for the Defendants submitted that it had been raised in the Defendants' opening skeleton argument, but I cannot see that it was. Since the issue is purely one of construction of the claim, however, I cannot see that there is any objection to the Defendants advancing the argument in closing submissions; it does not affect the expert evidence. I therefore turn to consider the argument on its merits.
242. Counsel for the Defendants sought to draw an analogy with Conversant's approach to integer A of the claim: on Conversant's construction of integer A, it was immaterial that the claim is silent as to what happens to scheduled transmissions. Likewise, he argued, it was immaterial that the claim is silent as to what happens if it is determined that the MAC is transmitting.
243. Counsel for Conversant submitted as follows. As discussed in paragraph 162 above, integer C is binary in nature, returning either a positive or negative result. Integer D specifies what is to happen if the result is negative. It is implicit that something different must happen if the result is positive, otherwise the check at integer C would be redundant. Reading the claim as a whole, it is also implicit that, if the result is positive, the MAC simply carries on transmitting for the remainder of the vTTI. Furthermore, that is consistent with the disclosure of the Patent at [0052].
244. Counsel for Conversant also pointed out that, as noted in paragraph 72 above, the Patent acknowledges Samsung. He submitted that it would not be right to construe claim 1 as almost reading onto Samsung.
245. Given that the argument is one of obviousness not anticipation, I am not convinced by the latter point, but nevertheless I accept Conversant's submissions as to how the claim should be construed. I do not consider that Conversant is being inconsistent in its approach to the claim, as counsel for the Defendants suggested.

Terry

246. *The disclosure.* Terry is directed to discontinuous transmissions (DTX) on the downlink in the context of a time division duplex (TDD) channel in UMTS. Given when it was written and its contents, the skilled reader would understand that it concerned an early UMTS release (probably Release 4). The skilled reader would therefore understand that the invention of Terry was designed to be implemented in a system that pre-dated high speed transmissions and the associated new functionality introduced in Releases 5 and 6, such as the 2 ms TTI, and changes to the layer 2 architecture in both UTRAN and the UE.

247. Terry concerns the Coded Composite Transport Channel (CCTrCH), a data channel which may contain one or more transport channels (TrCHs). Each TTI, a Transport Format Combination Index (TFCI) is sent on a control channel, relating to data carried on the CCTrCH. Full DTX mode (with which Terry is concerned) arises when, for a given CCTrCH, no data is provided by the upper protocol layers, in which case the Node B will not send any downlink transmissions on that CCTrCH for a given TTI. The start of full DTX is indicated by transmitting a Special Burst (SB) – a zero-valued TFCI. Thereafter SBs are sent periodically in accordance with the Special Burst Scheduling Period (SBSP) as long as the Node B remains in full DTX mode. The periodic bursts enable the UE to determine that the channel remains active, but that there is no data to transmit.
248. The problem to which Terry is directed is that the UE does not “know” the duration of the SBSP, so is required to remain ready to receive data (or an SB) on a channel for which full DTX mode has been initiated, even though no data may be being sent. That uses up processing power at the UE, which wastes battery life (an important issue in relation to a UE).
249. The solution which Terry proposes is to enable discontinuous reception at the UE i.e. to enable the UE to turn off processing during full DTX, when neither an SB nor data are being sent. In order to do so, however, the UE must therefore work out when it can expect to receive either the SB or new data, to avoid the risk either that new data or an SB will be transmitted while the UE is in a non-receiving state. This is achieved by coordinating the start of downlink transmissions with the SBSP period, coupled with the UE determining the SBSP, so that the UE can turn off its receiver periodically without missing the arrival of data or a further SB. Terry is, therefore, directed to improvements in discontinuous reception on the downlink of the CCTrCH channel, to improve UE battery life.
250. The detailed description in Terry therefore has essentially two parts: first, it describes methods for limiting the recommencement of downlink transmissions, once full DTX is initiated, to a period corresponding with the SBSP. This is described in [0017] – [0024] by reference to Figures 1 and 2. In summary, when the UTRAN MAC determines that there is no data to transmit in a given TTI, an idle period is initiated for the period of the SBSP, during which the UTRAN MAC does not process data, which is buffered in the RLC. At the end of the SBSP period, the UTRAN checks whether there is buffered data to transmit, and if not the process begins again. If there is data to be transmitted, then it is transmitted and the SBSP is then re-triggered. The SBSP is not a cyclic window, but a timer which is started each time the UTRAN MAC has no further data to send. Figure 2 depicts a slight variation on this method wherein data with certain quality of service parameters is sent as soon as it arrives in the RLC buffer, regardless of whether an idle period has been initiated.
251. The second part of Terry deals with methods for the UE to “learn” the SBSP period, to ensure that the periods where the UE receiver is turned off are coordinated with transmissions of the SB or data. This is the focus of the embodiments described in [0024] - [0048] by reference to Figures 3-6.
252. Figure 3 illustrates a process by which the UE is able to learn the SBSP, essentially by measuring the time between successive SBs. This is done iteratively, since one may not always get an accurate result, and Terry suggests at [0027] that five equivalent

values should be measured before the SBSP can be considered to have been determined.

253. Figure 4 illustrates a further method, once the SBSP has been determined, for the UE to assess whether the Node B is using the method of Terry (i.e. only sending new data in accordance with the SBSP). Here the UE monitors transmissions for new data, to see whether they are being sent in accordance with the SBSP or not – if they are then a counter is incremented, if not the counter is decremented, and once it reaches a suitable threshold the UE assumes that the Node B is using intelligent scheduling. It is only at that time that the UE is able to switch off its receiver between SBs.
254. The overall effect of these schemes is that the UE listens every TTI for potentially a very long time: at least for 10 full DTX cycles and possibly more. That could potentially represent 25 seconds, an age in telecommunications terms.
255. Prof Marshall’s evidence was that there are a number of aspects of the teaching in Terry that the skilled addressee would find unsatisfactory or puzzling. In cross-examination, he opined that the skilled person would simply “file” the document (i.e. in the bin). Dr Brydon accepted that the ideas disclosed by Terry were “not great”, and had nothing to say to the suggestion that they were terrible. So far as the evidence goes, Terry has never been implemented.
256. *Obviousness.* In approaching the Defendants’ case on obviousness over Terry, the starting point is that the Defendants do not allege that claim 1 of the Patent is obvious over common general knowledge.
257. The case which was put to Prof Marshall in cross-examination, however, was ultimately based on “what you would do once you have thrown Terry away”. This case was that the skilled person would appreciate from their common general knowledge that a simpler way to proceed was to have discontinuous transmission in the downlink. This would require the receiver to know when to listen for transmissions, and the skilled person would know from their common general knowledge that this could be done by having a cyclic series of defined starting points for transmissions. Furthermore, the skilled person would appreciate that they could take the same approach to the uplink. Accordingly, so it is said, Uplink DRX is obvious; and if Uplink DRX is within claim 1, then claim 1 is obvious.
258. In my judgment this will not do. If the claimed invention is not obvious over the common general knowledge, it cannot be obvious for the skilled person who has read Terry to throw Terry away and arrive at the claimed invention by applying common general knowledge concepts. Counsel for the Defendants also sought to rely upon evidence given by Dr Brydon that Terry disclosed a “high level idea” as well as detailed implementations of that idea, but the “high level idea” was the common general knowledge concept described above, and so the result is the same. At best, this is an approach based on hindsight.

Insufficiency

259. The Defendants have pleaded certain insufficiency objections as if they were squeezes on construction. As counsel for Conversant pointed out, however, at least the first two assume that the Defendants are right on points of construction. No coherent argument

was presented by the Defendants that the claims are insufficient if construed in the manner that I have concluded is correct.

Summary of principal conclusions

260. For the reasons given above, I conclude that:

- i) at least claim 1 of the Patent is essential to Uplink DRX, and therefore the Defendants have infringed at least claim 1 if it is valid; but
- ii) the granted claims are invalid for added matter, and the proposed amendment does not cure the invalidity, on the fourth ground advanced by the Defendants (but not on any of the other grounds); although
- iii) the claims are not obvious over either Samsung or Terry; and
- iv) the claims are not insufficient.